arXiv:1610.06506v1 [math.OC] 20 Oct 2016

ASTRO-DF: A CLASS OF ADAPTIVE SAMPLING TRUST-REGION ALGORITHMS FOR DERIVATIVE-FREE STOCHASTIC OPTIMIZATION

SARA SHASHAANI *, FATEMEH S. HASHEMI †, AND RAGHU PASUPATHY ‡

Abstract. We consider unconstrained optimization problems where only "stochastic" estimates of the objective function are observable as replicates from a Monte Carlo oracle. The Monte Carlo oracle is assumed to provide no direct observations of the function gradient. We present ASTRO-DF — a class of derivative-free trust-region algorithms, where a stochastic local interpolation model is constructed, optimized, and updated iteratively. Function estimation and model construction within ASTRO-DF is *adaptive* in the sense that the extent of Monte Carlo sampling is determined by continuously monitoring and balancing metrics of sampling error (or variance) and structural error (or model bias) within ASTRO-DF. Such balancing of errors is designed to ensure that Monte Carlo effort within ASTRO-DF is sensitive to algorithm trajectory, sampling more whenever an iterate is inferred to be close to a critical point and less when far away. We demonstrate the almost-sure convergence of ASTRO-DF's iterates to a first-order critical point when using linear or quadratic stochastic interpolation models. The question of using more complicated models, e.g., regression or stochastic kriging, in combination with adaptive sampling is worth further investigation and will benefit from the methods of proof presented here. We speculate that ASTRO-DF's iterates achieve the canonical Monte Carlo convergence rate, although a proof remains elusive.

Key words. derivative-free optimization, simulation optimization, stochastic optimization, trust-region

1. INTRODUCTION.

We consider unconstrained stochastic optimization (SO) problems, that is, optimization problems in continuous space where the objective function(s) can only be expressed implicitly via a Monte Carlo oracle. The Monte Carlo oracle is assumed to not provide any direct observations of the function derivatives.

SO has recently gathered attention due to its versatile formulation, allowing the user to specify functions involved in an optimization problem implicitly, e.g., through a stochastic simulation. As a result, SO allows virtually any level of problem complexity to be embedded, albeit at the possible price of a computationally burdensome and slow Monte Carlo oracle. SO has seen wide recent adoption — see, for example, applications in telecommunication networks [31], traffic control [42], epidemic forecasting [40] and health care [1]. Recent editions of the Winter Simulation Conference (`www.informs-sim.org`) have dedicated an entire track to the SO problem and its various flavors. For a library of SO problems, see `www.simopt.org` and [46, 47].

1.1. Problem Statement.

The SO problem we consider is formally stated as follows:

$$\text{Problem } P : \text{minimize } f(\boldsymbol{x}) \text{ subject to } \boldsymbol{x} \in \mathbb{R}^d, \tag{1.1}$$

where $f : \mathbb{R}^d \to \mathbb{R}$ is bounded from below and has Lipschitz continuous gradients. Furthermore, the function $f(\boldsymbol{x}) = \mathbb{E}[F(\boldsymbol{x})]$ is the expectation of a random function $F(\boldsymbol{x})$ that is observable through Monte Carlo. This means, for instance, that one can generate n identically distributed samples or replicates $F_i(\boldsymbol{x}), j = 1, 2, \ldots, n$ of $F(\boldsymbol{x})$ by "executing" the Monte Carlo simulation n times at the point $\boldsymbol{x}$. This leads to the estimator $\bar{F}(\boldsymbol{x}, n) = n^{-1} \sum_{j=1}^{n} F_j(\boldsymbol{x})$ having standard error estimated

*School of Industrial Engineering, Purdue University (email: sshashaa@purdue.edu)
†Department of Industrial and Systems Engineering, Virginia Tech (email: fatemeh@vt.edu)
‡Department of Statistics, Purdue University (email: pasupath@purdue.edu)

as $\hat{\sigma}_F(\boldsymbol{x},n)/\sqrt{n}$ where $\hat{\sigma}_F^2(\boldsymbol{x},n) = n^{-1}\sum_{j=1}^{n}\left(F_j(\boldsymbol{x}) - \bar{F}(\boldsymbol{x},n)\right)^2$. We assume that no direct observations of the gradient $\nabla f(\cdot)$, e.g., through IPA [4, p. 214], are available through the Monte Carlo oracle. This means that methods seeking a gradient estimate need to resort to indirect methods such as as finite differencing [4, p. 209], leading to biased estimators.

An algorithm for solving the above problem will be evaluated based on its ability to return a (random) sequence of iterates $\{\boldsymbol{X}_k\}$ converging in some rigorously defined probabilistic metric to a first- or second-order critical point of the function f. Thus, each "run" of a solution algorithm will return a random sequence of iterates $\{\boldsymbol{X}_k\}$, and SO algorithms that return iterate sequences $\{\boldsymbol{X}_k\}$ guaranteed to converge to a critical point with probability one will be called *consistent*.

1.2. Complications. The presence of a Monte Carlo oracle lends flexibility to the SO problem formulation, but it also brings with it a simply-stated complication: the lack of uniform deterministic error guarantees. Specifically, suppose $f(\boldsymbol{x},n)$ is the Monte Carlo estimate of the unknown desired function value $f(\boldsymbol{x})$ at the point $\boldsymbol{x}$, and n represents the extent of Monte Carlo effort. Then, simple probability arguments reveal that deterministic guarantees of the sort $|f(\boldsymbol{x},n) - f(\boldsymbol{x})| \leq \epsilon, \epsilon > 0$ do not hold irrespective of the size of n; instead, one has to be content with probabilistic precision guarantees of the form $\mathbb{P}\{|f(\boldsymbol{x},n) - f(\boldsymbol{x})| > \epsilon\} \leq \alpha$ for $n \geq n_0(\alpha)$. The analogous situation for function derivative estimation using Monte Carlo is worse. If the derivative estimate $\hat{\nabla} f(\boldsymbol{x}) := (\hat{\nabla}_1 f(\boldsymbol{x}), \hat{\nabla}_2 f(\boldsymbol{x}), \ldots, \hat{\nabla}_q f(\boldsymbol{x}))$ is constructed using a central-difference approximation as

$$\hat{\nabla}_i f(\boldsymbol{x}) = (2c_n)^{-1}(f(\boldsymbol{x} + c_n\boldsymbol{e}_i, n) - f(\boldsymbol{x} - c_n\boldsymbol{e}_i, n)), i = 1, 2, \ldots, q,$$

then, as in the function estimation context, no uniform guarantees on the accuracy of $\hat{\nabla} f(\boldsymbol{x})$ are available in general. Furthermore, the rate at which $\hat{\nabla} f(\boldsymbol{x})$ converges to $\nabla f(\boldsymbol{x})$ depends crucially on the choice of $\{c_n\}$, with the best possible rate $\mathcal{O}(n^{-1/3})$ under generic Monte Carlo sampling being much slower than the corresponding $\mathcal{O}(n^{-1/2})$ rate for function estimation. (See [4] for this and related results.) Most importantly, implementing such finite-difference derivative estimates within an SO algorithm is well recognized to be a delicate issue, easily causing instabilities. In any event, the lack of uniform deterministic guarantees in the SO context means that estimation error inevitably accumulates across iterations of an algorithm, thereby threatening convergence guarantees of the resulting iterates. Algorithms for solving SO have to somehow contend with such potential non-convergence due to mischance, either through the introduction of gain sequences as in stochastic approximation [34] or through appropriate sampling as in sample average approximation or retrospective approximation [33, 43].

A second complication within SO, but one that it partially shares with black-box deterministic optimization contexts, is the lack of information about function structure. Structural properties such as convexity, uni-modality, and differentiability, if known to be present, can be exploited when designing optimization algorithms. Such properties, when appropriate, are usually assumed within the deterministic context, and an appropriate solution algorithm devised. In SO, however, structural assumptions about the underlying true objective and constraint function, even if correct, may not provide as much leverage during algorithm development. This is because, due to the presence of stochastic error, the true objective and constraint functions are never directly observed; and, making structural assumptions about their observed

sample-paths is far more suspect.

REMARK 1. *Another aspect that is unique to SO is noteworthy. Monte Carlo calls are typically the most compute-intensive operations within SO contexts. And, depending on the nature of the SO algorithm, different number of Monte Carlo calls may be expended across iterations, e.g., constant in Stochastic Approximation (SA) [34], varying but predetermined in Retrospective Approximation (RA) [43], or random in Sampling Controlled Stochastic Recursion (SCSR) [45]. This means that the elemental measure of effort in SO — the number of Monte Carlo oracle calls — may not have a simple relationship with the notion of "iterations" defined within the specific SO algorithm, forcing a need for more careful book-keeping. This is why iterative SO algorithms are well-advised to measure convergence and convergence rates not in terms of the number of iterations, but rather in terms of the total number of Monte Carlo calls.*

1.3. ASTRO-DF and Overview of Contribution. Our particular focus in this research is that of developing a class of algorithms for solving low to moderate dimensional SO problems that have no readily discernible structure. We are inspired by the analogous problem in the deterministic context that has spurred the development of a special and arguably very useful class of optimization methods called model-based trust-region derivative-free (TRO-DF) algorithms [22, 49, 21, 5]. TRO-DF algorithms are typified by two aspects: (i) they eschew the direct computation and use of derivatives for searching, and instead rely on constructed models of guaranteed accuracy in specified "trust-regions"; (ii) the algorithmic search evolves by repeatedly constructing and optimizing a local model within a dynamic trust-region, explicitly restricting the distance between the successive iterates returned by the algorithm. The aspect in (i) is particularly suited for adaptation to SO contexts where direct derivative estimation can be delicate and unstable, requiring careful choice of step-sizes [4]; the aspect in (i) also aids efficiency because models constructed in previous iterations can be re-used with some updating, and no effort is expended for explicit estimation of derivatives. The aspect in (ii) runs counter to efficiency, but is designed to reduce variance in the algorithm's iterates, through steps that are more circumspect.

We construct a family of adaptive sampling trust-region optimization derivative-free (ASTRO-DF) algorithms for the SO context. In their most rudimentary form, ASTRO-DF algorithms follow a familiar idea for iteratively estimating the first and second order critical points of a function. Given a current random iterate $\boldsymbol{X}_k$ that approximates the first-order critical point of interest, ASTRO-DF constructs a tractable "local" stochastic model using Monte Carlo observations of the objective function at carefully chosen points around $\boldsymbol{X}_k$. The constructed model is then optimized within the local region in which it is constructed to obtain a candidate solution $\tilde{\boldsymbol{X}}_{k+1}$. Next, the objective function is *observed* (using Monte Carlo) at $\tilde{\boldsymbol{X}}_{k+1}$ and compared against the value *predicted* by the model at $\tilde{\boldsymbol{X}}_{k+1}$. If the observed decrease in function values from $\boldsymbol{X}_k$ to $\tilde{\boldsymbol{X}}_{k+1}$ exceeds the decrease predicted by the constructed model in a certain stochastic sense, the candidate $\tilde{\boldsymbol{X}}_{k+1}$ is accepted as the next iterate $\boldsymbol{X}_{k+1}$. As a vote of confidence on the constructed model, the trust-region radius is then expanded by a factor. Otherwise, that is, if the predicted decrease is much lower than the observed decrease (again, in a certain precise sense), the candidate $\tilde{\boldsymbol{X}}_{k+1}$ is rejected, the trust-region radius is shrunk, and the local model is updated in an attempt to improve accuracy. This iterative process then repeats to produce a random sequence of iterates $\{\boldsymbol{X}_k\}$ that is realized in each run of ASTRO-DF.

REMARK 2. *Throughout this paper, we use the term "sampling" to refer to the*

act of obtaining replicates using multiple runs of the Monte Carlo oracle at a fixed point. This is not to be confused with sampling design points in the search region. So, when we say that the sample size is n, we mean that n amount of Monte Carlo effort was expended to obtain the function estimate at a fixed point.

The above ideas for model construction, trust-region management, and candidate point acceptance say nothing about how much Monte Carlo effort to expend. Since all observations for function estimation and model construction are based on Monte Carlo, the resulting accuracy estimates are at best probabilistic, leading us to the question of how much to sample. Too little Monte Carlo effort threatens convergence due to accumulated stochastic and deterministic errors, and too much Monte Carlo sampling means reduced overall efficiency. Identifying the correct Monte Carlo sampling trade-off is more than a theoretical question, and answering it adequately entails more than broad prescriptions on sampling rates. To produce good implementations, sampling prescriptions ought to be automatic and specific to the problem at hand, which usually means relying on inference based on algorithm trajectory.

To resolve the issue of how much to sample, we propose that a simple strategy called *adaptive sampling* be incorporated within derivative-free trust-region algorithms in the SO context. Recognizing that the error in function and model estimation can be decomposed orthogonally into error due to sampling (or variance in the case of unbiased estimates) and error due to structure (or bias), adaptive sampling seeks to ensure that "just adequate" Monte Carlo sampling is performed by balancing these errors. For example, when constructing a local model, Monte Carlo sampling in ASTRO-DF is adaptive in the sense that sampling continues until a certain continuously monitored metric of model quality exceeds a metric of sampling variability. A similar rule is employed when estimating the objective function at a point for purposes of candidate acceptance. We believe that such adaptive sampling paves the way for efficiency because it reacts to the observed algorithm trajectory and, as we shall see, keeps the different sources of error within the algorithm in lock-step. The resulting algorithm remains practical because of the simplicity of the proposed adaptive sampling rule — sample until the estimated standard error falls below a certain specified power of the prevailing trust-region radius.

Remark 3. *Adaptive sampling as an idea is not new and has been used with great success in other areas such as sequential confidence interval construction [20, 28] and SO on finite spaces [30].*

Adaptive sampling, while invaluable as an implementation idea, introduces substantial complications when analyzing algorithm behavior. Akin to what happens during sequential sampling in the context of confidence interval construction [20, 28], the explicit dependence of the extent of Monte Carlo sampling on algorithm trajectory causes systematic early stopping and consequent bias in the function estimates obtained within ASTRO-DF. In other words, when using adaptive sampling, $\mathbb{E}[f(\boldsymbol{x}, n)] \neq f(\boldsymbol{x})$ in general since the sample size n is a stopping time [25, pp. 182] that will depend on $f(\boldsymbol{x}, n)$. Demonstrating that ASTRO-DF's iterates converge to a first-order critical point with probability one then entails demonstrating that the bias effects of adaptive sampling, especially when used within the derivative-free trust-region context, wear away asymptotically. We accomplish this by first generically characterizing a relationship between the moments of the adaptive sample size and the function estimates at stopping, and then showing that the errors induced due to model construction, algorithm recursion, and function estimation remain in lock-step (or comparable) throughout ASTRO-DF's evolution.

We note that ASTRO-DF, as presented here, assumes that a stochastic linear or stochastic quadratic interpolation model is constructed during the model-construction step of the algorithm. While such models are reasonable and have seen wide use in the analogous deterministic context, other possibly more powerful model construction techniques such as regression or stochastic kriging [3] should be considered in place of interpolation models, especially alongside adaptive sampling. Ongoing research investigates this question and it is our belief that the proof techniques that we present in this paper will carry over, albeit with some changes.

2. PRELIMINARIES. In this section, we list notation, key definitions, and some basic results that will be used invoked throughout the rest of the document.

2.1. Notation and Convention. We use bold font for vectors, script font for sets, lower case font for real numbers and upper case font for random variables. Hence $\{\boldsymbol{X}_k\}$ denotes a sequence of random vectors in $\mathbb{R}^d$, $\boldsymbol{x} = \left(x^1, x^2, \ldots, x^d\right)$ denotes a d-dimensional vector of real numbers, $\mathcal{Y} := \{\boldsymbol{Y}_1, \boldsymbol{Y}_2, \ldots, \boldsymbol{Y}_p\}$ denotes a set of p real vectors, and $\mathcal{Y} := \{\boldsymbol{Y}_1, \boldsymbol{Y}_2, \ldots, \boldsymbol{Y}_p\}$ denotes a set of p random vectors. The set $\mathcal{B}(\boldsymbol{x}; r) = \left\{\boldsymbol{y} \in \mathbb{R}^d : \ \|\boldsymbol{y} - \boldsymbol{x}\| \leq r\right\}$ is the closed ball of radius $r > 0$ with center $\boldsymbol{x}$.

For a sequence of random vectors $\{\boldsymbol{X}_k\}$, $\boldsymbol{X}_k \xrightarrow{wp1} \boldsymbol{X}$ denotes convergence with probability one or almost-sure convergence. For a sequence of real numbers $\{a_k\}$, we say $a_k = o(1)$ if $\lim_{k\to\infty} a_k = 0$; we say $a_k = \mathcal{O}(1)$ if $\{a_k\}$ is bounded, that is, there exists a constant $M > 0$ such that $|a_k| < M$ for large enough k. For sequences of real numbers $\{a_k\}$, $\{b_k\}$, we say that $a_k \sim b_k$ if $\lim_{k\to\infty} a_k/b_k = 1$. For a sequence of random variables $\{\boldsymbol{X}_k\}$, we say $\boldsymbol{X}_k = \mathcal{O}_p(1)$ if $\{\boldsymbol{X}_k\}$ is stochastically bounded, that is, given $\epsilon > 0$ there exists $M(\epsilon) \in \mathbb{R}$ such that $\mathbb{P}\{\boldsymbol{X}_k \in (-M(\epsilon), M(\epsilon))\} \geq 1 - \epsilon$ for all $k \geq K(\epsilon) \in \mathbb{N}$. For a sequence of sets $\{A_n\}$ defined on a probability space, the set $\mathbb{P}\{A_n \text{ i.o.}\} := \mathbb{P}\{\bigcap_{n=1}^{\infty} \bigcup_{m=n}^{\infty} A_m\}$ refers to the event that "A_n happens infinitely often."

2.2. Key Definitions. The following definitions will be invoked heavily during our exposition and analysis of ASTRO-DF. For further details on these definitions, consult [22] and [41].

DEFINITION 2.1. *(Poised and Λ-Poised Sets) Given $\boldsymbol{x} \in \mathbb{R}^d$ and $\Delta > 0$, let $\mathcal{Y} = \{\boldsymbol{Y}_i \in \mathcal{B}(\boldsymbol{x}; \Delta),\ i = 1, 2, \ldots, p\}$ be a finite set and $\Phi(\boldsymbol{z}) = \left(\phi^1(\boldsymbol{z}), \phi^2(\boldsymbol{z}), \ldots, \phi^q(\boldsymbol{z})\right)$ be a polynomial basis on $\mathbb{R}^d$. Define*

$$P(\Phi, \mathcal{Y}) = \begin{bmatrix} \phi^1(\boldsymbol{Y}_1) & \phi^2(\boldsymbol{Y}_1) & \ldots & \phi^q(\boldsymbol{Y}_1) \\ \phi^1(\boldsymbol{Y}_2) & \phi^2(\boldsymbol{Y}_2) & \ldots & \phi^q(\boldsymbol{Y}_2) \\ \vdots & \vdots & \vdots & \vdots \\ \phi^1(\boldsymbol{Y}_p) & \phi^2(\boldsymbol{Y}_p) & \ldots & \phi^q(\boldsymbol{Y}_p) \end{bmatrix}. \tag{2.1}$$

Then, $\mathcal{Y}$ is said to be a "poised set" in $\mathcal{B}(\boldsymbol{x}; \Delta)$ if the matrix $P(\Phi, \mathcal{Y})$ is nonsingular. A poised set $\mathcal{Y}$ is said to be "Λ-poised" in $\mathcal{B}(\boldsymbol{x}; \Delta)$ if

$$\Lambda \geq \max_{j=1,\ldots,p} \max_{\boldsymbol{z} \in \mathcal{B}(\boldsymbol{x}; \Delta)} |\ell_j(\boldsymbol{z})|,$$

where $\ell_j(\boldsymbol{z})$ are the Lagrange polynomials associated with $\mathcal{Y}$.

DEFINITION 2.2. *(Polynomial Interpolation Models) Let $f : \mathbb{R}^d \subseteq \mathbb{R}^d \to \mathbb{R}$ be a real-valued function and let $\mathcal{Y}$ and Φ be as defined in Definition 2.1 with $p = q$. Suppose we can find $\boldsymbol{\alpha} = \left(\alpha^1, \alpha^2, \ldots, \alpha^p\right)$ such that*

$$P(\Phi, \mathcal{Y})\, \boldsymbol{\alpha} = \left(f(\boldsymbol{Y}_1), \ldots, f(\boldsymbol{Y}_p)\right)^T. \tag{2.2}$$

(Such an $\boldsymbol{\alpha}$ is guaranteed to exist if $\mathcal{Y}$ is poised.) Then the function $m(\boldsymbol{z}) : \mathcal{B}(\boldsymbol{x};\Delta) \to \mathbb{R}$ given by

$$m(\boldsymbol{z}) = \sum_{j=1}^{p} \alpha^j \phi^j(\boldsymbol{z}) \tag{2.3}$$

is said to be a polynomial interpolation model of f on $\mathcal{B}(\boldsymbol{x};\Delta)$. As a special case, $m(\boldsymbol{z})$ is said to be a linear interpolation model of f on $\mathcal{B}(\boldsymbol{x};\Delta)$ if $\Phi(\boldsymbol{z}) := (\phi^1, \phi^2, \ldots, \phi^p) = (1, z^1, z^2, \ldots, z^d)$, and a quadratic interpolation model of f on $\mathcal{B}(\boldsymbol{x};\Delta)$ if $\Phi(\boldsymbol{z}) := (\phi^1, \phi^2, \ldots, \phi^p) = \left(1, z^1, z^2, \ldots, z^d, \frac{1}{2}(z^1)^2, z^1 z^2, \ldots, \frac{1}{2}(z^2)^2, \ldots, \frac{1}{2}(z_d)^2\right)$.

Definition 2.3. *(Stochastic Interpolation Models) A model constructed as in Definition 2.2 but with sampled function estimates is called a stochastic interpolation model. Specifically, analogous to (2.2), suppose $\hat{\boldsymbol{\alpha}} = \left(\hat{\alpha}^1, \hat{\alpha}^2, \ldots, \hat{\alpha}^p\right)$ is such that*

$$P(\Phi, \mathcal{Y})\,\hat{\boldsymbol{\alpha}} = \left(\bar{F}(\boldsymbol{Y}_1, n(\boldsymbol{Y}_1)), \bar{F}(\boldsymbol{Y}_2, n(\boldsymbol{Y}_2)), \ldots, \bar{F}(\boldsymbol{Y}_p, n(\boldsymbol{Y}_p))\right)^T,$$

Then the stochastic function $M(\boldsymbol{z}) : \mathcal{B}(\boldsymbol{x};\Delta) \to \mathbb{R}$ given as $M(\boldsymbol{z}) = \sum_{j=1}^{p} \hat{\alpha}^j \phi^j(\boldsymbol{z})$ is said to be a stochastic polynomial interpolation model of f on $\mathcal{B}(\boldsymbol{x};\Delta)$, where $\Phi(\boldsymbol{z}) = \left(\phi^1(\boldsymbol{z}), \phi^2(\boldsymbol{z}), \ldots, \phi^q(\boldsymbol{z})\right)$ and $P(\Phi, \mathcal{Y})$ are as in Definition 2.2.

Definition 2.4. *(Fully-linear and Fully-quadratic Models) Given $\boldsymbol{x} \in \mathbb{R}^d$, $m(\boldsymbol{z}) : \mathcal{B}(\boldsymbol{x};\Delta) \to \mathbb{R}$, $m \in \mathcal{C}^1$ is said to be a $(\kappa_{ef}, \kappa_{eg})$-fully-linear model of f on $\mathcal{B}(\boldsymbol{x};\Delta)$ if it has a Lipschitz continuous gradient with Lipschitz constant ν_{gL}^m, and there exist constants κ_{ef}, κ_{eg} (not dependent on $\boldsymbol{z}$ and Δ) such that*

$$\begin{aligned} |f(\boldsymbol{z}) - m(\boldsymbol{z})| &\le \kappa_{ef}\Delta^2; \\ \|\nabla f(\boldsymbol{z}) - \nabla m(\boldsymbol{z})\| &\le \kappa_{eg}\Delta. \end{aligned} \tag{2.4}$$

Similarly $m(\boldsymbol{z}) : \mathcal{B}(\boldsymbol{x};\Delta) \to \mathbb{R}$, $m \in \mathcal{C}^2$ is said to be a $(\kappa_{ef}, \kappa_{eg}, \kappa_{eh})$-fully-quadratic model of f on $\mathcal{B}(\boldsymbol{x};\Delta)$ if it has a Lipschitz continuous second derivative with Lipschitz constant ν_{hL}^m, and constants κ_{ef}, κ_{eg}, κ_{eh} (not dependent on $\boldsymbol{z}$, Δ) such that

$$\begin{aligned} |f(\boldsymbol{z}) - m(\boldsymbol{z})| &\le \kappa_{ef}\Delta^3; \\ \|\nabla f(\boldsymbol{z}) - \nabla m(\boldsymbol{z})\| &\le \kappa_{eg}\Delta^2; \\ \left\|\nabla^2 f(\boldsymbol{z}) - \nabla^2 m(\boldsymbol{z})\right\| &\le \kappa_{eH}\Delta. \end{aligned} \tag{2.5}$$

A linear interpolation model constructed using a poised set $\mathcal{Y}$ can be shown to be $(\kappa_{ef}, \kappa_{eg})$-fully-linear; likewise, a quadratic interpolation model constructed using a poised set $\mathcal{Y}$ can be shown to be $(\kappa_{ef}, \kappa_{eg}, \kappa_{eh})$-fully-quadratic.

Definition 2.5. *(Cauchy Reduction) Step $\boldsymbol{s}$ is said to achieve κ_{fcd} fraction of Cauchy reduction for $m(\cdot)$ on $\mathcal{B}(\boldsymbol{x};\Delta)$ with some $\Delta > 0$, if*

$$m(\boldsymbol{x}) - m(\boldsymbol{x}+\boldsymbol{s}) \ge \frac{\kappa_{fcd}}{2}\|\nabla m(\boldsymbol{x})\| \min\left\{\frac{\|\nabla m(\boldsymbol{x})\|}{\|\nabla^2 m(\boldsymbol{x})\|}, \Delta\right\}, \tag{2.6}$$

where $\nabla m(\boldsymbol{x})$ and $\nabla^2 m(\boldsymbol{x})$ are the model gradient and the model Hessian at point $\boldsymbol{x}$. We assume $\|\nabla m(\boldsymbol{x})\| / \left\|\nabla^2 m(\boldsymbol{x})\right\| = +\infty$ when $\nabla^2 m(\boldsymbol{x}) = \boldsymbol{0}$. A Cauchy step with $\kappa_{fcd} = 1$ is obtained by minimizing the model $m(\cdot)$ along the steepest descent direction within $\mathcal{B}(\boldsymbol{x};\Delta)$ [22, p. 175]. Accordingly, the Cauchy step is especially easy to obtain when $m(\cdot)$ is linear or quadratic.

2.3. Useful Results. We will now state some basic results that will be used at various points in the paper. The first of these (Theorem 2.6) is a seminal result that is routinely used in sequential sampling contexts, especially when constructing confidence intervals. The second result (Theorem 2.7) is a variation of Lemma 2 and Theorem 1 in [27] which, again, originally appeared in the context of sequential confidence intervals. As we shall see, we will use Theorem 2.7 extensively in our analysis, to analyze the behavior of estimators that are constructed using sequential sampling. The postulates and the setting of Theorem 2.7 differ only a little from the setting of Theorem 1 in [27]; for this reason, we have chosen not to include a proof.

THEOREM 2.6 (Chow and Robbins, 1965). *Suppose random variables* $X_i, i = 1, 2, \ldots$ *are iid with variance* $\sigma^2 < \infty$, $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i$, $\hat{\sigma}_n^2 = n^{-1}\sum_{i=1}^n \left(X_i - \bar{X}_n\right)^2$, *and* $\{a_n\}$ *a sequence of positive constants such that* $a_n \to a$ *as* $n \to \infty$. *If*

$$N(d) = \inf\left\{n \geq 1 : \frac{\hat{\sigma}_n}{\sqrt{n}} \leq \frac{d}{a_n}\right\},$$

then $d^2 N(d)/\left(a^2\sigma^2\right) \xrightarrow{wp1} 1$ *and* $\hat{\sigma}_N/\sigma \xrightarrow{wp1} 1$ *as* $d \to 0$.

THEOREM 2.7. *Suppose random variables* $X_i, i = 1, 2, \ldots$ *are iid with* $\mathbb{E}[X_1] = 0, \mathbb{E}[X_1^2] = \sigma^2 > 0$, *and* $\mathbb{E}[|X_1|^{4v}] < \infty$ *for some* $v \geq 2$. *Let* $\hat{\sigma}_n^2 = n^{-1}\sum_{i=1}^n \left(X_i - \bar{X}_n\right)^2$, *where* $\bar{X}_n = n^{-1}\sum_{i=1}^n X_i$. *If*

$$N(\lambda) = \inf\left\{n \geq \lambda^\gamma : \frac{\hat{\sigma}_n}{\sqrt{n}} \leq \frac{\kappa}{\sqrt{\lambda}}\right\}, \gamma \in (0, 1]$$

then the following hold.

(i) As $\lambda \to \infty$,

$$\mathbb{P}\{N(\lambda) < \infty\} = 1 \text{ and } N(\lambda) \xrightarrow{wp1} \infty.$$

(ii) As $\lambda \to \infty$ *and for every* $s < v$,

$$\mathbb{E}[N^s(\lambda)] \sim \sigma^{2s}\kappa^{-2s}\lambda^s.$$

(iii) For every $\epsilon \in (0, 1)$,

$$\mathbb{P}\{N(\lambda) \leq \sigma^2\kappa^{-2}\lambda(1-\epsilon)\} = \beta\lambda^{-(v-1)\gamma}.$$

where, as [27] notes, β *is a generic constant that might depend only on* v *and the moments of* X_1 *but not on* λ.

(iv) As $\lambda \to \infty$,

$$\mathbb{E}[\bar{X}_{N(\lambda)}^2(\lambda)] \sim \kappa^2\lambda^{-1}.$$

We next state Lemma 2.8 which characterizes the error in the stochastic interpolation model introduced in Definition 2.3. Lemma 2.8 is essentially a stochastic variant of a result that appears in Chapter 3 of [22]. We provide a sketch of the proof of Lemma 2.8 in the Appendix.

LEMMA 2.8. *Let* $\mathcal{Y} = \{\boldsymbol{Y}_1, \boldsymbol{Y}_2, \ldots, \boldsymbol{Y}_p\}$ *be a* Λ*-poised set on* $\mathcal{B}(\boldsymbol{Y}_1; \Delta)$. *Let* $m(\boldsymbol{z})$ *be an interpolation model of* f *on* $\mathcal{B}(\boldsymbol{Y}_1; \Delta)$. *Let* $M(\boldsymbol{z})$ *be the corresponding stochastic interpolation model of* f *on* $\mathcal{B}(\boldsymbol{Y}_1; \Delta)$ *constructed using observations* $\bar{F}(\boldsymbol{Y}_i, n(\boldsymbol{Y}_i)) = f(\boldsymbol{Y}_i) + E_i$ *for* $i = 1, 2, \ldots, p$.

(i) For all $\boldsymbol{z} \in \mathcal{B}(\boldsymbol{Y}_1; \Delta)$,

$$|M(\boldsymbol{z}) - m(\boldsymbol{z})| \leq p\Lambda \max_{i=1,2,\ldots,p} \left|\bar{F}(\boldsymbol{Y}_i, n(\boldsymbol{Y}_i)) - f(\boldsymbol{Y}_i)\right|.$$

(ii) If $M(\boldsymbol{z})$ *is a stochastic linear interpolation model of* f *on* $\mathcal{B}(\boldsymbol{Y}_1; \Delta)$*, then there exist positive constants* $\kappa_{egL1}, \kappa_{egL2}$ *such that for* $\boldsymbol{z} \in \mathcal{B}(\boldsymbol{Y}_1; \Delta)$,

$$\|\nabla M(\boldsymbol{z}) - \nabla f(\boldsymbol{z})\| \leq \kappa_{egL1}\Delta + \kappa_{egL2}\frac{\sqrt{\sum_{i=2}^{d+1}(E_i - E_1)^2}}{\Delta}.$$

If $M(\boldsymbol{z})$ *is a stochastic quadratic interpolation model of* f *on* $\mathcal{B}(\boldsymbol{Y}_1; \Delta)$*, then there exist positive constants* $\kappa_{egQ1}, \kappa_{egQ2}$ *such that*

$$\|\nabla M(\boldsymbol{z}) - \nabla f(\boldsymbol{z})\| \leq \kappa_{egQ1}\Delta^2 + \kappa_{egQ2}\frac{\sqrt{\sum_{i=2}^{(d+1)(d+2)/2}(E_i - E_1)^2}}{\Delta}.$$

We end this section by stating two basic results that we repeatedly invoke, and which can be found in most standard treatments of probability such as [12]. The first of these is used to upper bound the probability of the union of events; the second provides sufficient conditions to ensure that an infinite number of a collection $\{A_n\}$ of events happening is zero.

LEMMA 2.9 (Boole's Inequality). *Let* $A_1, A_2, \cdots$ *be a countable set of events defined on a probability space. Then* $\mathbb{P}(\bigcup_i A_i) \leq \sum_i \mathbb{P}(A_i)$. *Particularly, we see that if the random variables* $X, X_i,\ i = 1, 2, \cdots, q$ *satisfy* $X \leq X_1 + X_2 + \cdots + X_q$*, then*

$$\begin{aligned}(X > c) &\subseteq (X_1 + X_2 + \cdots X_q > c)\\ &\subseteq \left(X_1 > \frac{c}{q}\right) \cup \left(X_2 > \frac{c}{q}\right) \cup \cdots \cup \left(X_q > \frac{c}{q}\right),\end{aligned}$$

implying (from Boole's inequality) that

$$\mathbb{P}\{X > c\} \leq \mathbb{P}\left\{\bigcup_{i=1}^{q}\left(X_i > \frac{c}{q}\right)\right\} \leq \sum_{i=1}^{q}\mathbb{P}\left\{X_i > \frac{c}{q}\right\}.$$

LEMMA 2.10 (Borel-Cantelli's First Lemma). *For a sequence* $A_1, A_2, \ldots$ *of events defined on a probability space, if* $\sum_{n=1}^{\infty}\mathbb{P}\{A_n\} < \infty$*, then the probability* $\mathbb{P}\{A_n \text{ i.o.}\}$ *of* A_n *happening "infinitely often" is*

$$\mathbb{P}\{A_n \text{ i.o.}\} := \mathbb{P}\left\{\bigcap_{n=1}^{\infty}\bigcup_{m=n}^{\infty} A_m\right\} = 0.$$

3. RELATED WORK. Much progress has been made in recent times on solving various flavors of the SO problem. The predominant solution methods in the simulation literature fall into two broad categories called Stochastic Approximation (SA) and Sample-Average Approximation (SAA). SA and SAA have enjoyed a long history with mature theoretical and algorithmic literature. More recently, newer classes of algorithms that can be described as "stochastic versions" of iterative structures in the deterministic context have emerged.

3.1. SA and SAA. Virtually all stochastic approximation type methods are subsumed by the following generic form:

$$\boldsymbol{X}_{k+1} = \Pi_{\mathcal{D}}\left(\boldsymbol{X}_k - a_k \boldsymbol{G}_k\right), \tag{3.1}$$

where $\Pi_{\mathcal{D}}(\boldsymbol{x})$ is the projection of the point $\boldsymbol{x}$ onto the set $\mathcal{D}$, $\{a_k\}$ is a user-chosen positive-valued scalar sequence, and $\boldsymbol{G}_k$ is an estimator of the gradient $\nabla f(\boldsymbol{X}_k)$ of the function f at the point $\boldsymbol{X}_k$. When direct Monte Carlo observations of the objective function f are available, the most common expression for $\boldsymbol{G}_k = \left(G_k^1, G_k^2, \cdots, G_k^d\right)$ is either the central-difference approximation $G_k^i = (2c_k)^{-1}(F(\boldsymbol{X}_k + c_k\boldsymbol{e}_i) - F(\boldsymbol{X}_k - c_k\boldsymbol{e}_i))$ or the forward-difference approximation $G_k^i = c_k^{-1}(F(\boldsymbol{X}_k + c_k\boldsymbol{e}_i) - F(\boldsymbol{X}_k))$ of the gradient $\nabla f(\boldsymbol{X}_k)$, where $\{c_k\}$ is a positive-valued sequence and $F : \mathbb{R}^d \to \mathbb{R}$ is the observable estimator of the objective function $f : \mathbb{R}^d \to \mathbb{R}$. The resulting recursion is the famous Kiefer-Wolfowitz process [32, 15]. More recent recursions include an estimated Hessian $H_k(\cdot)$ of the function f at the point $\boldsymbol{X}_k$:

$$\boldsymbol{X}_{k+1} = \Pi_{\mathcal{D}}\left(\boldsymbol{X}_k - a_k H_k^{-1}\boldsymbol{G}_k\right), \tag{3.2}$$

making the resulting recursion in (3.1) look closer to the classical Newton's iteration in the deterministic context. The Hessian estimator $H_k(\cdot)$ has d^2 entries, and hence, most methods that use (3.2) estimate $H_k(\cdot)$ either using a parsimonious design (e.g., [54, 55]), or construct it from the history of observed points.

As can be seen in (3.1), the SA recursion is simply stated and implemented, and little has changed in its basic structure since 1951, when it was first introduced by Robbins and Monro [50] for the context of finding a zero of a "noisy" vector function. Instead, much of the research over the ensuing decades has focused on questions such as convergence and convergence rates of SA type algorithms, the effect of averaging on the consistency and convergence rates of the iterates, and efforts to choose the sequence $\{a_k\}$ in an adaptive fashion. Some good entry points into the vast SA literature include [34, 38, 48, 44]. See [17, 16, 57] recent attempts to address the persistent dilemma of choosing the gain sequence $\{a_k\}$ to ensure good practical performance.

SAA, in contrast to SA, is more a framework than an algorithm to solve SO problems. Instead of solving Problem P, SAA asks to solve a "sample-path" Problem P_n (to optimality) to obtain a solution estimator $\boldsymbol{X}_n$. Formally, in the unconstrained context, SAA seeks to solve

$$\text{Problem } P_n: \text{ minimize } f(\boldsymbol{x}, n) \text{ subject to } \boldsymbol{x} \in \mathbb{R}^d, \tag{3.3}$$

where $f(\boldsymbol{x}, n)$ is computed using a "fixed" sample of size n.

SAA is attractive in that Problem P_n becomes a *deterministic* optimization problem and SAA can bring to bear all of the advances in deterministic nonlinear programming methods [11, 39] of the last few decades. SAA has been the subject of a tremendous amount of theoretical and empirical research over the last two decades. For example, the conditions that allow the transfer of structural properties from the sample-path to the limit function $f(\boldsymbol{x})$ [33, Propositions 1,3,4]; the sufficient conditions for the consistency of the optimal value and solution of Problem P_n assuming the numerical procedure in use within SAA can produce global optima [53, Theorem 5.3]; consistency of the set of stationary points of Problem P_n [53, 6]; convergence rates for the optimal value [53, Theorem 5.7] and optimal solution [33, Theorem 12]; expressions for the minimum sample size m that provides probabilistic guarantees on the

optimality gap of the sample-path solution [52, Theorem 5.18]; methods for estimating the accuracy of an obtained solution [37, 8, 9]; and quantifications of the trade-off between searching and sampling [51], have all been thoroughly studied. SAA is usually not implemented in the vanilla form P_n due to known issues relating to an appropriate choice of the sample size n. There have been recent advances [44, 24, 8, 9, 10] aimed at defeating the issue of sample size choice.

3.2. Stochastic TRO. Two algorithms that are particularly noteworthy competitors to what we propose here are STORM [19] and the recently proposed algorithm by Larson and Billups [35] (henceforth LB2014). While the underlying logic in both of these algorithms key differences arise in terms of what has been assumed about the quality of the constructed models and how such quality can be achieved in practice. Another notable difference is that STORM also treats the context of biased estimators, that is, contexts where $\mathbb{E}[f(\boldsymbol{x}, n)] \neq f(\boldsymbol{x})$. A key postulate that guarantees consistency in STORM is that the constructed models are of a certain specified quality (characterized through the notion of probabilistic full linearity) with a probability exceeding a fixed threshold. The authors provide a way to construct such models using function estimates constructed as sample means. Crucially, the prescribed sample means in STORM use a sample size that is derived using the Chebyshev inequality with an assumed upper bound on the variance. By contrast, the sample sizes in ASTRO-DF are determined adaptively by balancing squared bias and variance estimates for the function estimator. While this makes the sample size in ASTRO-DF a stopping time [12] thereby complicating proofs, such adaptive sampling enables ASTRO-DF to differentially sample across the search space, leading to efficiency.

LB2014, like STORM, uses random models. Unlike STORM, however, the sequence of models constructed in LB2014 are assumed to be accurate (as measured by a certain rigorous notion) with a probability sequence that converges to one. A related version of LB2014 [13] addresses the case of differing levels of (spatial) stochastic error through the use of weighted regression schemes, where the weights are chosen heuristically.

Another noteworthy algorithm for the context we consider in this paper is VNSP, proposed by Deng and Ferris [23, 24]. VNSP uses a quadratic interpolation model within a trust-region optimization framework, and is derivative-free in the sense that only function estimates are assumed to be available. Model construction, inference, and improvement, along with (nondecreasing) sample size updates happen within a Bayesian framework with an assumed Gaussian conjugate prior. Convergence theory for VNSP is accordingly within a Bayesian setting.

In the slightly more tangential context where unbiased gradient estimates are assumed to be available, a number of trust-region type algorithms have emerged in the last decade or so. STRONG or Stochastic Trust-Region Response-Surface Method [18], for instance, is an adaptive sampling trust-region algorithm for solving SO problems that is in the spirit of what we propose here. A key feature of STRONG is local model construction through a design of experiments combined with a hypothesis testing procedure. STRONG assumes that the error in the derivative observations are additive and have a Gaussian distribution. Amos et. al. [2] and Bastin et. al. [7] are two other examples of algorithms that treat the setting where unbiased observations of the gradient are assumed to be available. (The former, in fact, assumes that unbiased estimates of the Hessian of the objective function are available.) Bastin et. al. [7] is specific to the problem of estimation within mixed-logit models.

Algorithm 1 ASTRO-DF Main Algorithm

Require: Initial guess $\boldsymbol{x}_0 \in \mathbb{R}^d$, initial trust-region radius $\tilde{\Delta}_0 > 0$ and maximum radius $\Delta_{\max} > 0$, model "fitness" threshold $\eta_1 > 0$, trust-region expansion constant $\gamma_1 > 1$ and contraction constant $\gamma_2 \in (0,1)$, initial sample size n_0, sample size lower bound sequence $\{\lambda_k\}$ such that $k^{(1+\epsilon)} = \mathcal{O}(\lambda_k)$, initial sample set $\tilde{\mathcal{Y}}_0 = \{\boldsymbol{x}_0\}$, and outer adaptive sampling constant κ_{oas}.

1: **for** $k = 0, 1, 2, \ldots$ **do**

2: *Model Construction:* Construct the model at $\boldsymbol{X}_k$ by calling Algorithm 2 with the candidate trust-region radius $\tilde{\Delta}_k$ and candidate set of sample points $\tilde{\mathcal{Y}}_k$,

$$[M_k(\boldsymbol{X}_k + \boldsymbol{s}), \Delta_k, \mathcal{Y}_k] = \texttt{AdaptiveModelConstruction}(\tilde{\Delta}_k, \tilde{\mathcal{Y}}_k).$$

Set $\widetilde{N}_k = N(\boldsymbol{X}_k)$.

3: *TR Subproblem:* Approximate the kth step by minimizing the model in the trust-region, $\boldsymbol{S}_k = \operatorname{argmin}_{\|\boldsymbol{s}\| \le \Delta_k} M_k(\boldsymbol{X}_k + \boldsymbol{s})$, and set the new candidate point $\tilde{\boldsymbol{X}}_{k+1} = \boldsymbol{X}_k + \boldsymbol{S}_k$.

4: *Evaluate:* Estimate the function at the candidate point using adaptive sampling to obtain $\bar{F}(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1})$, where

$$\widetilde{N}_{k+1} = \max\left\{\lambda_k, \min\left\{n : \frac{\hat{\sigma}_F\left(\tilde{\boldsymbol{X}}_{k+1}, n\right)}{\sqrt{n}} \le \frac{\kappa_{oas}\Delta_k^2}{\sqrt{\lambda_k}}\right\}\right\}, \tag{4.1}$$

Update:

5: Compute the success ratio $\hat{\rho}_k$ as

$$\hat{\rho}_k = \frac{\bar{F}\left(\boldsymbol{X}_k, \widetilde{N}_k\right) - \bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right)}{M_k(\boldsymbol{X}_k) - M_k(\tilde{\boldsymbol{X}}_{k+1})}.$$

6: **if** $\hat{\rho}_k > \eta_1$ **then**

7: $\boldsymbol{X}_{k+1} = \tilde{\boldsymbol{X}}_{k+1}$, $\tilde{\Delta}_{k+1} = \min\{\gamma_1\Delta_k, \Delta_{\max}\}$, $N_{k+1} = \widetilde{N}_{k+1}$. Update the sample set $\tilde{\mathcal{Y}}_{k+1}$ to include the new iterate.

8: **else**

9: $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k$, $\tilde{\Delta}_{k+1} = \gamma_2\Delta_k$, $N_{k+1} = \widetilde{N}_k$. Update the sample set $\tilde{\mathcal{Y}}_{k+1}$, if needed, to include the rejected candidate point.

10: **end if**

11: **end for**

4. ASTRO–DF OVERVIEW AND ALGORITHM LISTING. ASTRO-DF is an adaptive sampling trust-region derivative-free algorithm whose essence is encapsulated within four repeating stages: (i) local stochastic model construction and certification through adaptive sampling; (ii) constrained optimization of the constructed model for identifying the next candidate solution; (iii) re-estimation of the objective function at the next candidate solution through adaptive sampling; and (iv) iterate and trust-region update based on a (stochastic) sufficient decrease check. These stages appear with italic labels in Algorithm 1. In what follows, we describe each step of Algorithm 1 in further detail.

In Step 2 of Algorithm 1, a stochastic model of the function $f(\cdot)$ in the trust-region $\mathcal{B}(\boldsymbol{X}_k; \Delta_k)$ is constructed using Algorithm 2. The aim of Algorithm 2 is to construct a model of a specified quality within a trust-region having radius smaller than a fixed multiple of the model gradient norm. During the j_kth iteration of Algorithm 2, a poised set $\mathcal{Y}_k^{(j_k)} \triangleq \{\boldsymbol{Y}_1^{(j_k)}, \boldsymbol{Y}_2^{(j_k)}, \ldots, \boldsymbol{Y}_p^{(j_k)}\}$ in the "candidate" trust-region having radius $\tilde{\Delta}_k w^{j_k - 1}$ and center $\boldsymbol{Y}_1^{(j_k)} = \boldsymbol{X}_k$ is chosen (Step 3); Monte Carlo function estimates are then obtained at each of the points in $\mathcal{Y}_k^{(j_k)}$ with $N\left(\boldsymbol{Y}_i^{(j_k)}\right)$ being the sample size at point $\boldsymbol{Y}_i^{(j_k)}$ after the j_kth iteration of the contraction loop. Sampling at each point in $\mathcal{Y}_k^{(j_k)}$ is adaptive and continues (Steps 4–6) until the estimated standard errors

Algorithm 2 $[M_k(\boldsymbol{X}_k+\boldsymbol{s}), \Delta_k, \mathcal{Y}_k]$=`AdaptiveModelConstruction`$(\tilde{\Delta}_k, \tilde{\mathcal{Y}}_k)$

Require: *Parameters from ASTRO-DF:* candidate trust-region radius $\tilde{\Delta}_k$ and candidate sample set $\tilde{\mathcal{Y}}_k$ (possibly with cardinality $< p$).
Parameters specific to `AdaptiveModelConstruction:` trust-region contraction factor $w \in (0,1)$, trust-region and gradient balance constant μ, gradient inflation constant β with $0 < \beta < \mu$, and inner adaptive sampling constant κ_{ias}.

1: Initialize $j_k = 1$, set $\mathcal{Y}_k^{(j_k)} = \tilde{\mathcal{Y}}_k$, and set $\boldsymbol{Y}_1 = \boldsymbol{X}_k$ where $\boldsymbol{X}_k$ is the first element of $\tilde{\mathcal{Y}}_k$.
Contraction loop:
2: **repeat**
3: Improve $\mathcal{Y}_k^{(j_k)} = \left\{\boldsymbol{Y}_1^{(j_k)}, \boldsymbol{Y}_2^{(j_k)}, \ldots, \boldsymbol{Y}_p^{(j_k)}\right\}$ by appropriately choosing $\boldsymbol{Y}_i^{(j_k)},\ i = 2, 3, \cdots, p$ to make it a poised set in $\mathcal{B}(\boldsymbol{X}_k; \tilde{\Delta}_k w^{j_k-1})$.
4: **for** $i = 1$ to p **do**
5: Estimate $\bar{F}\left(\boldsymbol{Y}_i^{(j_k)}, N\left(\boldsymbol{Y}_i^{(j_k)}\right)\right)$, where

$$N\left(\boldsymbol{Y}_i^{(j_k)}\right) = \max\left\{\lambda_k, \min\left\{n : \frac{\hat{\sigma}_F\left(\boldsymbol{Y}_i, n\right)}{\sqrt{n}} \leq \frac{\kappa_{ias}(\tilde{\Delta}_k w^{j_k-1})^2}{\sqrt{\lambda_k}}\right\}\right\}. \tag{4.2}$$

6: **end for**
7: Construct a quadratic model $M_k^{(j_k)}\left(\boldsymbol{X}_k+\boldsymbol{s}\right)$ via interpolation.
8: Set $j_k = j_k + 1$.
9: **until** $\tilde{\Delta}_k w^{j_k-1} \leq \mu \| \nabla M_k^{(j_k)}\left(\boldsymbol{X}_k\right) \|$.
10: Set $M_k(\boldsymbol{X}_k+\boldsymbol{s}) = M_k^{(j_k)}(\boldsymbol{X}_k+\boldsymbol{s})$, $\nabla M_k(\boldsymbol{X}_k) = \nabla M_k^{(j_k)}(\boldsymbol{X}_k)$, and $\nabla^2 M_k(\boldsymbol{X}_k) = \nabla^2 M_k^{(j_k)}(\boldsymbol{X}_k)$.
11: **return** $M_k(\boldsymbol{X}_k+\boldsymbol{s})$, $\Delta_k = \min\left\{\tilde{\Delta}_k, \max\left\{\beta \left\|\nabla M_k(\boldsymbol{X}_k)\right\|, \tilde{\Delta}_k w^{j_k-1}\right\}\right\}$, and $\mathcal{Y}_k = \mathcal{Y}_k^{(j_k)}$.

$\hat{\sigma}_F\left(\boldsymbol{Y}_i^{(j_k)}, N\left(\boldsymbol{Y}_i^{(j_k)}\right)\right) / \sqrt{N\left(\boldsymbol{Y}_i^{(j_k)}\right)}$ of the function estimates $\bar{F}\left(\boldsymbol{Y}_i^{(j_k)}, N\left(\boldsymbol{Y}_i^{(j_k)}\right)\right)$ drop below a slightly inflated square of the candidate trust-region radius. A linear (or quadratic) interpolation model is then constructed using the obtained function estimates in Step 5. (If a linear interpolation model is constructed, $p = d+1$, and if a quadratic interpolation model is constructed, $p = (d+1)(d+2)/2$.) If the resulting model $M_k^{(j_k)}(\boldsymbol{z}), \boldsymbol{z} \in \mathcal{B}(\boldsymbol{X}_k; \tilde{\Delta}_k w^{j_k-1})$ is such that the candidate trust-region radius $\tilde{\Delta}_k w^{j_k-1}$ is too large compared to the norm of the model gradient $\left\|\nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\|$, that is, if $\tilde{\Delta}_k w^{j_k-1} > \mu \left\|\nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\|$, then the candidate trust-region radius is shrunk by a factor w and control is returned back to Step 3. On the other hand, if the candidate trust-region radius is smaller than the product of μ and the norm of the model gradient, then the resulting stochastic model is accepted but over an updated incumbent trust-region radius given by Step 11. (Step 11 of Algorithm 2, akin to [21], updates the incumbent trust-region radius to the point in the interval $[\tilde{\Delta}_k w^{j_k-1}, \tilde{\Delta}_k]$ that is closest to $\beta \left\|\nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\|$).

We emphasize the following four issues pertaining to Step 2 in Algorithm 1 and the model resulting from the application of Algorithm 2.

(i) Due to the nature of the chosen poised set $\mathcal{Y}_k$, the (hypothetical) limiting model $m_k(\boldsymbol{X}_k)$ constructed from true function observations on $\mathcal{Y}_k$ will be $(\kappa_{ef}, \kappa_{eg})$-fully-linear (or $(\kappa_{ef}, \kappa_{eg}, \kappa_{eh})$-fully-quadratic) on the updated trust-region $\mathcal{B}(\boldsymbol{X}_k; \Delta_k)$. Of course, the model $m_k(\boldsymbol{X}_k)$ is unavailable since true function evaluations are unavailable; and it makes no sense to talk about whether the constructed model $M_k(\boldsymbol{X}_k)$ is $(\kappa_{ef}, \kappa_{eg})$-fully-linear (or $(\kappa_{ef}, \kappa_{eg}, \kappa_{eh})$-fully quadratic) since it is constructed from stochastic function estimates.

(ii) By construction, the trust-region resulting from the application of Algorithm 2 has a radius that is at most β times the model gradient norm $\|\nabla M_k(\boldsymbol{X}_k)\|$.

(iii) The structure of adaptive sampling in Step 5 of Algorithm 2 is identical to that appearing for estimation in Step 4 of Algorithm 1. The adaptive sampling step simply involves sampling until the estimated standard error of the function estimate comes within a factor of the deflated square of the incumbent trust-region radius. As our convergence proofs will reveal, balancing the estimated standard error to any lower power of the incumbent trust-region radius will threaten consistency of ASTRO-DF's iterates.

(iv) As can be seen from the algorithm listings, the quality of the constructed model in ASTRO-DF is ensured (in Step 2) at every iteration. As will be evident from our analysis, such stringency eases theoretical analysis. This is in contrast to many modern implementations of derivative-free trust-region algorithms where model quality is checked and improved (if necessary) through a "criticality step" which is triggered only when the norm of the model gradient falls below a threshold. This selective model quality assurance makes the algorithm more numerically efficient, at the price of a more complicated trust-region management and convergence analysis (See acceptable iterations and model improving iterations in [22, p. 185]). We believe that incorporating a similar criticality step in ASTRO-DF will ease computational burden during implementation.

Let us now resume our discussion of Algorithm 1. In Step 2, Algorithm 1 executes `AdaptiveModelConstruction` to obtain a model $M_k(\boldsymbol{z}), \boldsymbol{z} \in \mathcal{B}(\boldsymbol{X}_k; \Delta_k)$ whose limiting approximation is $(\kappa_{ef}, \kappa_{eg})$-fully-linear (or $(\kappa_{ef}, \kappa_{eg}, \kappa_{eh})$-fully-quadratic) as observed in (i) above. Step 3 in Algorithm 1 then approximately solves the constrained optimization problem $\boldsymbol{S}_k = \arg\min_{\|\boldsymbol{s}\| \le \Delta_k} M_k(\boldsymbol{X}_k + \boldsymbol{s})$ to obtain a candidate point $\tilde{\boldsymbol{X}}_{k+1} = \boldsymbol{X}_k + \boldsymbol{S}_k$ satisfying the κ_{fcd}-Cauchy decrease as defined in Assumption 4.

In preparation for checking if the candidate solution $\tilde{\boldsymbol{X}}_{k+1}$ provides sufficient decrease, Step 4 of Algorithm 1 obtains Monte Carlo samples of the objective function at $\tilde{\boldsymbol{X}}_{k+1}$, until the estimated standard error $\hat{\sigma}_F\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) / \sqrt{\widetilde{N}_{k+1}}$ of $\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right)$ is smaller than a slightly deflated square of the trust-region radius $\lambda_k^{-1/2} \kappa_{oas} \Delta_k^2$, subject to the sample size being at least as big as λ_k (see Remark 4).

In Step 5 of Algorithm 1, the obtained function estimate is used to check if the so called *success ratio* $\hat{\rho}_k$, that is, the ratio of the predicted to the observed function decrease at the point $\tilde{\boldsymbol{X}}_{k+1}$, exceeds a fixed threshold η_1. The denominator of the success ratio $\hat{\rho}_k$ is calculated by evaluating the constructed model at $\tilde{\boldsymbol{X}}_{k+1}$ using the analytical form listed in Definition 2.3. If $\hat{\rho}_k$ exceeds the threshold η_1, the candidate $\tilde{\boldsymbol{X}}_{k+1}$ is accepted as the new iterate $\boldsymbol{X}_{k+1}$, the iteration is deemed successful, and the trust-region is expanded (Step 6). If $\hat{\rho}_k$ falls below the specified threshold η_1, the candidate $\tilde{\boldsymbol{X}}_{k+1}$ is rejected (though it may remain in the sample set), the iteration is deemed unsuccessful, and the trust-region is shrunk (Step 9). In either case, $\tilde{\Delta}_{k+1}$ is set as the incumbent trust-region radius, N_{k+1} is set as the current sample size of $\boldsymbol{X}_{k+1}$, and $\mathcal{Y}_k$ is set as the interpolation set for the next iteration. Note that in the next iteration the sample size of $\boldsymbol{X}_{k+1}$ is subject to change through Step 2 again.

REMARK 4. *The sequence $\{\lambda_k\}$ appearing as the first argument of the "max" function in the expression for the adaptive sample size (in Step 4 of Algorithm 1 and Step 5 of Algorithm 2) is standard for all adaptive sampling contexts, e.g., [20, 28], and intended to nullify the effects of mischance without explicitly participating in the*

limit. Since $N = \max\left\{\lambda_k, \lambda_k \hat{\sigma}_F^2(\boldsymbol{X}_k, N)\kappa_{oas}^{-2}\Delta_k^{-4}\right\}$*, and since we will demonstrate that* $\hat{\sigma}_F^2(\boldsymbol{X}_k, N) \xrightarrow{wp1} \sigma > 0$ *and* $\Delta_k \xrightarrow{wp1} 0$*, the probability of the first argument in the expression for the adaptive sample size being binding will decay to zero as* $k \to \infty$*.*

5. CONVERGENCE ANALYSIS OF ASTRO-DF. As noted in Section 4, the convergence behavior of ASTRO-DF depends crucially on the behavior of three error terms expressed through the following decomposition:

$$\begin{aligned}\left|\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) - M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \leq & \left|\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) - f\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \\ & + \left|f\left(\tilde{\boldsymbol{X}}_{k+1}\right) - m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \\ & + \left|m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right) - M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|. \end{aligned} \tag{5.1}$$

The three terms appearing on the right-hand side of (5.1) can be interepreted, respectively, as follows: (i) the *stochastic sampling error* $\left|\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) - f\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|$ arising due to the fact that function evaluations are estimated using Monte Carlo; (ii) the *deterministic model error* $\left|f\left(\tilde{\boldsymbol{X}}_{k+1}\right) - m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|$ arising due to the choice of local model; and (iii) the *stochastic interpolation error* $\left|m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right) - M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|$ arising due to the fact that model prediction at unobserved points is a combination of the model bias and the error in (i). (The analysis in the deterministic context involves only the error in (ii).) Accordingly, driving the errors in (i) and (ii) to zero sufficiently fast, while ensuring the fully-linear or quadratic sufficiency of the expected model, guarantees almost sure convergence.

Driving the errors in (i) and (ii) to zero sufficiently fast is accomplished by forcing the sample sizes to increase across iterations at a sufficiently fast rate, something that we ensure by keeping the estimated standard error of all function estimates in lock step with the square of the trust-region radius. The trust-region radius is in turn also kept in lock-step with the model gradient through the model construction Algorithm 2. Such a deliberate lock-step between the model error, trust-region radius, and the model gradient is aimed at efficiency without sacrificing consistency.

In what follows, we provide a formal proof of the wp1 convergence of ASTRO-DF's iterates. Recall that we assume that the models being constructed within Step 2 of Algorithm 1 are either linear or quadratic. Furthermore, we focus only on convergence to a first-order critical point of the function f.

We first list six standing assumptions that are assumed to hold for the ensuing results. Additional assumptions will be made as and when required.

ASSUMPTION 1. *The function* f *is continuously differentiable and bounded from below. Furthermore, it has Lipschitz continuous gradients, that is, there exists* ν_{gL} *such that* $\|\nabla f(\boldsymbol{x}) - \nabla f(\boldsymbol{y})\| \leq \nu_{gL}\|\boldsymbol{x} - \boldsymbol{y}\|$ *for all* $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^d$*.*

ASSUMPTION 2. *There exists a first-order critical point associated with Problem P, that is, there exists* $\boldsymbol{x}^* \in \mathbb{R}^d$ *such that:* $\nabla f(\boldsymbol{x}^*) = 0$*.*

ASSUMPTION 3. *The Monte Carlo oracle, when executed at* $\boldsymbol{X}_k$*, generates independent and identically distributed random variates* $F_j(\boldsymbol{X}_k) = f(\boldsymbol{X}_k) + \xi_j \,|\, \mathcal{F}_k$*, where* $\xi_1, \xi_2, \ldots$ *is a martingale-difference sequence adopted to* $\mathcal{F}_k$ *such that* $\mathbb{E}\left[\xi_j^2 \,|\, \mathcal{F}_k\right] = \sigma^2$ *for all* k*, where* $\sigma^2 < \infty$ *and* $\sup_k \mathbb{E}[|\xi_j|^{4v} \,|\, \mathcal{F}_k] < \infty$ *for some* $v \geq 2$*.*

Assumptions 1 – 3 are arguably mild assumptions relating to the problem. The following three assumptions relate to the nature of the proposed algorithm ASTRO-

DF.

ASSUMPTION 4. *The minimizer obtained in the trust-region subproblem (Step 3 of Algorithm 1) satisfies a κ_{fcd}-Cauchy decrease with $\kappa_{fcd} > 0$, that is,*

$$M_k\left(\boldsymbol{X}_k\right) - M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right) \geq \frac{\kappa_{fcd}}{2}\left\|\nabla M_k\left(\boldsymbol{X}_k\right)\right\| \min\left\{\frac{\left\|\nabla M_k\left(\boldsymbol{X}_k\right)\right\|}{\left\|\nabla^2 M_k\left(\boldsymbol{X}_k\right)\right\|}, \Delta_k\right\}.$$

ASSUMPTION 5. *There exists a positive constant κ_{bhm}, such that the model Hessian $\nabla^2 M_k\left(\boldsymbol{X}_k\right)$ satisfies*

$$\mathbb{P}\left\{\omega : \limsup_{k\to\infty}\left\|\nabla^2 M_k\left(\boldsymbol{X}_k(\omega)\right)\right\| \leq \kappa_{bhm}\right\} = 1,$$

that is, the model Hessians are eventually bounded for all $\boldsymbol{x}$ and k almost surely.

ASSUMPTION 6. *The "lower-bound sequence" $\{\lambda_k\}$ is chosen to satisfy $k^{(1+\epsilon)} = \mathcal{O}(\lambda_k)$ for some $\epsilon > 0$.*

Assumption 4 is usually easily ensured through an appropriate choice in the trust-region subproblem step (Step 3 of Algorithm 1). For example, Assumption 4 is satisfied if the optimization resulting from Step 3 of Algorithm 1 yields a solution that is at least as good as the Cauchy step t_C, which is the minimizer of the model $M_k(\cdot)$ along the direction $\nabla M_k(\boldsymbol{X}_k)$ and constrained to the trust-region, that is,

$$t_C = \underset{\alpha\in[0,\Delta_k]}{\operatorname{argmin}}\ M_k(\boldsymbol{X}_k - \alpha\nabla M_k(\boldsymbol{X}_k)).$$

(See Section 10.1 in [22] for additional details.) Likewise, Assumption 5 can be enforced through a check that is performed each time the model is constructed or updated. And, Assumption 6 imposes a (weak) minimum increase on the sample sizes for estimation and model construction operations within ASTRO-DF.

REMARK 5. *It is our view that the minimum rate of increase on the lower bound sequence $\{\lambda_k\}$ can be reduced to a logarithmic increase instead of what has been assumed in Assumption 6. Using the notation of Theorem 2.7, this will require a large-deviation type bound on the tail probability $\mathbb{P}\{|\bar{X}_N| > t\}$ after assuming the existence of the moment-generating function of X_i's. To the best of our knowledge there currently exist no such results for fixed-width confidence interval stopping, which is the context of Theorem 2.7.*

5.1. Main Results. We are now ready to establish the *consistency*, that is, the almost sure convergence to a stationary point, of the iterates generated by ASTRO-DF. The roadmap for consistency consists of three main theorems supported by three lemmas. The first of the main theorems is Theorem 5.2 where we demonstrate that the sequence of trust-region radii $\{\Delta_k\}$ across the iterations of ASTRO-DF converges to zero with probability one. Next, Theorem 5.4 establishes that the iterations within ASTRO-DF are eventually successful with probability one. The final result that establishes the almost sure convergence of ASTRO-DF's iterates to a stationary point appears as Theorem 5.6. With the exception of Lemma 5.5, all results that follow analyze the ensemble probabilistic behavior of the sample paths to then make inferences about individual sample-paths, primarily through the first Borel-Cantelli's first lemma (see Lemma 2.10). Lemma 5.5 is an exception in that the analysis there is pathwise, without probabilistic arguments.

We start with Lemma 5.1 which establishes that the sequence of function estimates at the iterates generated by ASTRO-DF has to remain bounded with probability one.

LEMMA 5.1. *Let Assumptions 1, 3, and 6 hold. Then*

(a) $\mathbb{P}\left\{\lim_{k\to\infty} \bar{F}\left(\boldsymbol{X}_k, N_k\right) = -\infty\right\} = 0$,
(b) $\mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right| \geq c\Delta_k^2 \text{ i.o.}\right\} = 0$.

Proof. For part (a) we know from Assumption 1 that f is bounded from below. Hence we can write for any $c \in \mathbb{R}$,

$$\mathbb{P}\left\{\lim_{k\to\infty} \bar{F}\left(\boldsymbol{X}_k, N_k\right) = -\infty\right\} \leq \mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right| \geq c \text{ i.o.}\right\}, \tag{5.2}$$

where i.o. stands for "infinitely often." (See Section 2.1 for a formal definition and use of "infinitely often.") However by the law of total probability,

$$\begin{aligned}\mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right| \geq c\right\} &= \mathbb{E}\left[\mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right| \geq c \,|\, \mathcal{F}_k\right\}\right] \\ &\leq \mathbb{E}[c^{-2}\mathbb{E}[\left(\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right)^2 \,|\, \mathcal{F}_k]], \end{aligned} \tag{5.3}$$

where the last inequality follows from Chebyshev's inequality [12]. Now invoke part (iv) of Theorem 2.7 along with Assumption 3 and the sample size expression in (4.1) to notice that

$$\mathbb{E}[\left(\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right)^2 \,|\, \mathcal{F}_k] \sim \kappa_{oas}^2 \Delta_k^4 \lambda_k^{-1} \tag{5.4}$$

as $\lambda_k \to \infty$; that is, for large-enough k, we can write for any $\delta > 0$ that

$$\begin{aligned}\mathbb{E}[\left(\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right)^2 \,|\, \mathcal{F}_k] &\leq (1+\delta)\kappa_{oas}^2 \Delta_k^4 \lambda_k^{-1} \\ &\leq (1+\delta)\kappa_{oas}^2 \Delta_{\max}^4 \lambda_k^{-1}. \end{aligned} \tag{5.5}$$

Now use (5.3) and (5.5) to write

$$\begin{aligned}\mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right| \geq c\right\} &\leq \mathbb{E}[c^{-2}\mathbb{E}[\left(\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right)^2 \,|\, \mathcal{F}_k]] \\ &\leq \mathbb{E}[c^{-2}(1+\delta)\kappa_{oas}^2 \Delta_{\max}^4 \lambda_k^{-1}] \\ &= c^{-2}(1+\delta)\kappa_{oas}^2 \Delta_{\max}^4 \lambda_k^{-1}. \end{aligned} \tag{5.6}$$

The right-hand side of (5.6) is summable since $k^{(1+\epsilon)} = \mathcal{O}(\lambda_k)$ for some $\epsilon > 0$; we can thus invoke the first Borel-Cantelli Lemma [12] and conclude that the right-hand side of (5.2) is zero. This proves part (a).

In fact, similar to (5.6) we have

$$\begin{aligned}\mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right| \geq c\Delta_k^2\right\} &\leq \mathbb{E}[c^{-2}\mathbb{E}[\left(\bar{F}\left(\boldsymbol{X}_k, N_k\right) - f\left(\boldsymbol{X}_k\right)\right)^2 \,|\, \mathcal{F}_k]] \\ &\leq \mathbb{E}[c^{-2}\Delta_k^{-4}(1+\delta)\kappa_{oas}^2 \Delta_k^4 \lambda_k^{-1}] \\ &= c^{-2}(1+\delta)\kappa_{oas}^2 \lambda_k^{-1}, \end{aligned} \tag{5.7}$$

which is again summable and hence allows the invocation of Borel-Cantelli's first lemma [12], proving part (b).
□

Next, we state a theorem that plays a crucial role in proving the overall convergence of ASTRO-DF iterates. Recall that even in deterministic derivative-free trust-region algorithms, unlike trust-region algorithms where derivative observations are available, the trust-region radius necessarily needs to decay to zero to ensure convergence. Theorem 5.2 states that this is indeed the case for ASTRO-DF. The proof

rests on Lemma 5.1 and the assumed sufficient Cauchy decrease guarantee during Step 3 of Algorithm 1.

THEOREM 5.2. *Let Assumptions 1, 3, 4, 5 and 6 hold. Then* $\Delta_k \xrightarrow{wp1} 0$ *as* $k \to \infty$.

Proof. We note during the ith iteration in Step 2 of Algorithm 1 that $\bar{F}(\boldsymbol{X}_i, N_i)$ and $\bar{F}\left(\boldsymbol{X}_i, \widetilde{N}_i\right)$ denote the function estimate at the point $\boldsymbol{X}_i$ *before* entering and upon exiting `AdaptiveModelConstruction` respectively.

$$\bar{F}(\boldsymbol{X}_k, N_k) = \bar{F}(\boldsymbol{X}_1, N_1) + \sum_{i=1}^{k-1} (A_i + B_i) \tag{5.8}$$

where the summands $A_i = \bar{F}(\boldsymbol{X}_{i+1}, N_{i+1}) - \bar{F}\left(\boldsymbol{X}_i, \widetilde{N}_i\right)$ and $B_i = \bar{F}\left(\boldsymbol{X}_i, \widetilde{N}_i\right) - \bar{F}(\boldsymbol{X}_i, N_i)$. In words A_i represents the reduction in the function estimates during the ith iteration and B_i represents the difference between the two estimates of the function at the point $\boldsymbol{X}_i$ at the end of iteration $i-1$ and i. We now make the following observations about A_i and B_i.

(a) If i is an unsuccessful iteration, then $A_i = 0$ since $\boldsymbol{X}_i = \boldsymbol{X}_{i+1}$.

(b) If i is a successful iteration, we know by definition that $\hat{\rho}_i \geq \eta_1$. If we denote $\kappa_{efd} = (2\mu)^{-1}\eta_1\kappa_{fcd}\min\left\{(\mu\kappa_{bhm})^{-1}, 1\right\}$, then by Assumptions 4 and 5, and by the assurance in Algorithm 2 that $\Delta_k \leq \mu\|\nabla M_k(\boldsymbol{X}_k)\|$, we have

$$\begin{aligned} A_i &\leq \eta_1 \left(M_i(\boldsymbol{X}_{i+1}) - M_i(\boldsymbol{X}_i)\right) \\ &\leq -\frac{\eta_1}{2}\kappa_{fcd}\|\nabla M_i(\boldsymbol{X}_i)\| \min\left\{\frac{\|\nabla M_i(\boldsymbol{X}_i)\|}{\|\nabla^2 M_i(\boldsymbol{X}_i)\|}, \Delta_i\right\} \\ &\leq -\kappa_{efd}\Delta_i^2. \end{aligned} \tag{5.9}$$

(c) For any given $c > 0$, Lemma 5.1 ensures that $\mathbb{P}\{|B_i| > c \text{ i.o.}\} = 0$ since

$$\begin{aligned} \mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_i, \widetilde{N}_i\right) - \bar{F}(\boldsymbol{X}_i, N_i)\right| > c\right\} &\leq \mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{X}_i, \widetilde{N}_i\right) - f(\boldsymbol{X}_i)\right| > \frac{c}{2}\right\} \\ &\quad + \mathbb{P}\left\{\left|f(\boldsymbol{X}_i) - \bar{F}(\boldsymbol{X}_i, N_i)\right| > \frac{c}{2}\right\}, \end{aligned}$$

using Boole's inequality (see Lemma 2.9). This implies that except for a set of measure zero, $|B_i| \leq c$ for large enough i.

Now suppose $\mathcal{D} := \{\omega : \lim_{k\to\infty}\Delta_k(\omega) \neq 0\}$ denotes the set of sample-paths for which the trust-region radius does not decay to zero. For contraposition, suppose $\mathcal{D}$ has positive measure. Consider a sample-path $\omega_0 \in \mathcal{D}$. Since unsuccessful iterations are necessarily contracting iterations, we can find $\delta(\omega_0) > 0$ and the sub-sequence of successful iterations $\{k_j\}$ in the sample-path ω_0 such that $\Delta_{k_j}(\omega_0) \geq \delta(\omega_0)$. This implies from observation (b) above that

$$A_{k_j}(\omega_0) \leq -\kappa_{efd}\delta^2(\omega_0). \tag{5.10}$$

Now, observation (a) above implies that

$$A_{k_j+\ell}(\omega_0) \leq 0, \ell = 1, 2, \ldots, k_{j+1} - k_j - 1. \tag{5.11}$$

Also by the observation (c) above, and choosing $c = \frac{1}{3}\kappa_{efd}\delta^2(\omega_0)$, we see that for large-enough i,

$$\left|\bar{F}\left(\boldsymbol{X}_i(\omega_0), \widetilde{N}_i(\omega_0)\right) - \bar{F}(\boldsymbol{X}_i(\omega_0), N_i(\omega_0))\right| \leq \frac{2}{3}\kappa_{efd}\delta^2(\omega_0). \tag{5.12}$$

We then write for large-enough j,

$$\begin{aligned}\sum_{\ell=k_j}^{k_{j+1}-1}\left(A_\ell(\omega_0)+B_\ell(\omega_0)\right) &= A_{k_j}(\omega_0)+\sum_{\ell=k_j}^{k_{j+1}-1}B_\ell(\omega_0)\\ &\le A_{k_j}(\omega_0)\\ &\quad+\bar{F}\left(\boldsymbol{X}_{k_{j+1}-1}(\omega_0),\widetilde{N}_{k_{j+1}-1}(\omega_0)\right)\\ &\quad-\bar{F}\left(\boldsymbol{X}_{k_j+1}(\omega_0),N_{k_j+1}(\omega_0)\right)\\ &\le -\frac{1}{3}\kappa_{efd}\delta^2(\omega_0),\end{aligned}\tag{5.13}$$

where the first equality follows from observation (a) above, the first inequality follows from the definition of B_ℓ, and the second inequality follows from (5.10) and (5.12). The inequality in (5.13) (and the fact that there is an entire sequence $\{k_j\}$ of successful iterations) means that $\lim_{k\to\infty}\bar{F}\left(\boldsymbol{X}_k\left(\omega_0\right),N_k(\omega_0)\right)=-\infty$ thus contradicting Lemma 5.1. The assertion of the theorem thus holds.□

Relying on Theorem 5.2, we now show that the model gradient converges to the true gradient almost surely. This, of course, does not imply that the true gradient itself converges to zero — a fact that will be established subsequently. Implicit in the proof of Theorem 5.2 is the requirement that Algorithm 2 terminates in finite time with probability one, a fact that we establish through Lemma C.1 in the Appendix.

LEMMA 5.3. *Let Assumptions 1, 3 – 6 hold. Then* $\|\nabla M_k\left(\boldsymbol{X}_k\right)-\nabla f\left(\boldsymbol{X}_k\right)\|\xrightarrow{wp1}0$ *as* $k\to\infty$.

Proof. In Step 3 of Algorithm 2, $\tilde{\Delta}_k w^{j_k-1}$ denotes the trust-region radius over which the model is constructed. (Note that due to Step 11 of Algorithm 2, $\tilde{\Delta}_k w^{j_k-1}$ may or may not equal the ending trust-region radius Δ_k upon completion of k iterations of ASTRO-DF.) Then, we know from part (ii) of Lemma 2.8 that

$$\|\nabla M_k\left(\boldsymbol{X}_k\right)-\nabla f\left(\boldsymbol{X}_k\right)\|\le\kappa_1\left(\tilde{\Delta}_k w^{j_k-1}\right)^\theta+\kappa_2\frac{\sqrt{\sum_{i=2}^p\left(E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right)^2}}{\left(\tilde{\Delta}_k w^{j_k-1}\right)},$$

where $E_{k,i}^{(j_k)}=\bar{F}\left(\boldsymbol{Y}_i^{(j_k)},N\left(\boldsymbol{Y}_i^{(j_k)}\right)\right)-f\left(\boldsymbol{Y}_i^{(j_k)}\right)$ for $i=1,\ldots,p$ denotes the error due to sampling at point $\boldsymbol{Y}_i^{(j_k)}$ after the j_kth iteration of the contraction loop. (Recall that $\boldsymbol{Y}_1^{(j_k)}=\boldsymbol{X}_k$; $p=d+1$ and $\theta=1$ in the linear interpolation models, and $p=(d+1)(d+2)/2$ and $\theta=2$ in the quadratic interpolation models. For the quantities κ_1 and κ_2 refer to part (ii) of Lemma 2.8.)

From Theorem 5.2, $\Delta_k\xrightarrow{wp1}0$ as $k\to\infty$, and hence, $\tilde{\Delta}_k w^{j_k-1}\xrightarrow{wp1}0$ as $k\to\infty$. Also, $\sqrt{\sum_{i=2}^p\left(E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right)^2}\le\sum_{i=2}^p\sqrt{\left(E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right)^2}=\sum_{i=2}^p\left|E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right|$. Considering these two observations, it suffices to show that as $k\to\infty$,

$$\left(\tilde{\Delta}_k w^{j_k-1}\right)^{-1}\sum_{i=2}^p\left|E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right|\xrightarrow{wp1}0.\tag{5.14}$$

Towards this, we write for $c>0$, large enough k, some $\delta>0$, and by Boole's inequality

(Lemma 2.9),

$$\begin{aligned}
\mathbb{P}\left\{\frac{\sum_{i=2}^{p}\left|E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right|}{\left(\tilde{\Delta}_k w^{j_k-1}\right)}\geq c\right\} &\leq \sum_{i=2}^{p}\mathbb{E}\left[\mathbb{P}\left\{\left|E_{k,i}^{(j_k)}-E_{k,1}^{(j_k)}\right|\geq \frac{c\left(\tilde{\Delta}_k w^{j_k-1}\right)}{p-1}\,|\,\mathcal{F}_k\right\}\right] \\
&\leq \sum_{i=2}^{p}\left(\mathbb{E}\left[\mathbb{P}\left\{\left|E_{k,i}^{(j_k)}\right|\geq \frac{c\left(\tilde{\Delta}_k w^{j_k-1}\right)}{2(p-1)}\,|\,\mathcal{F}_k\right\}\right]\right. \\
&\qquad \left.+\mathbb{E}\left[\mathbb{P}\left\{\left|E_{k,1}^{(j_k)}\right|\geq \frac{c\left(\tilde{\Delta}_k w^{j_k-1}\right)}{2(p-1)}\,|\,\mathcal{F}_k\right\}\right]\right) \\
&\leq 2(p-1)^3 4\left(c\tilde{\Delta}_k w^{j_k-1}\right)^{-2}(1+\delta)\kappa_{ias}^2(\tilde{\Delta}_k w^{j_k-1})^4\lambda_k^{-1} \\
&\leq 8(p-1)^3 c^{-2}(1+\delta)\kappa_{ias}^2\Delta_k^2\lambda_k^{-1}, \qquad (5.15)
\end{aligned}$$

where the second inequality above follows from the application of Boole's inequality (see Lemma 2.9), and the penultimate inequality above follows from arguments identical to those leading to (5.6) in the proof of Lemma 5.1 after using the adaptive sample size expression in (4.2). Since the right-hand side of (5.15) is summable, we can invoke the first Borel-Cantelli lemma [12] to conclude that (5.14) holds. □

We now show that for large enough iteration k, the steps within ASTRO-DF are always successful with probability one. This result is important in that it implies that the model gradient and the trust-region radius will remain in lock-step for large k, almost surely. The proof proceeds by dividing the model error into three components, each of which is shown to be controlled with probability one.

THEOREM 5.4. *Let Assumptions 1 – 6 hold. Then* $\mathbb{P}\{\hat{\rho}_k < \eta_1, \text{i.o.}\} = 0$ *for any* $\eta_1 \in (0,1)$.

Proof. At the end of Step 9 of Algorithm 2, let $m_k^{(j_k)}(\boldsymbol{z})$ be the interpolation model of f constructed on the poised set $\mathcal{Y}_k$. (Of course, we cannot construct $m_k(\cdot)$ explicitly because the true function values are unknown.) Then $m_k^{(j_k)}(\boldsymbol{z})$ is a $(\kappa_{ef},\kappa_{eg})$-fully-linear model of f on $\mathcal{B}\left(\boldsymbol{X}_k;\tilde{\Delta}_k w^{j_k-1}\right)$ and since $\Delta_k \geq \tilde{\Delta}_k w^{j_k-1}$, by the Lemma in [22, p. 200] we have that $m_k(\cdot)$ is a $(\kappa_{ef},\kappa_{eg})$-fully-linear model of f on $\mathcal{B}(\boldsymbol{X}_k;\Delta_k)$. In addition, Algorithm 2 ensures that $\Delta_k \leq \mu\|\nabla M_k(\boldsymbol{X}_k)\|$.

Assumption 4 on the Cauchy decrease in the minimization problem implies that

$$\begin{aligned}
M_k(\boldsymbol{X}_k)-M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right) &\geq \frac{\kappa_{fcd}}{2}\|\nabla M_k(\boldsymbol{X}_k)\|\min\left\{\frac{\|\nabla M_k(\boldsymbol{X}_k)\|}{\|\nabla^2 M_k(\boldsymbol{X}_k)\|},\Delta_k\right\} \\
&\geq \frac{\kappa_{fcd}}{2}\|\nabla M_k(\boldsymbol{X}_k)\|\min\left\{\frac{\Delta_k}{\mu\kappa_{bhm}},\Delta_k\right\} \\
&\geq \kappa_{md}\Delta_k^2. \qquad (5.16)
\end{aligned}$$

where $\kappa_{md} = (2\mu\kappa_{bhm})^{-1}\min(\mu\kappa_{bhm}^{-1},1)\kappa_{fcd}$. Recall that

$$\hat{\rho}_k := \frac{\bar{F}\left(\boldsymbol{X}_k,\widetilde{N}_k\right)-\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1},\widetilde{N}_{k+1}\right)}{M_k(\boldsymbol{X}_k)-M_k(\tilde{\boldsymbol{X}}_{k+1})}$$

and that $\bar{F}\left(\boldsymbol{X}_k,\widetilde{N}_k\right) = M_k(\boldsymbol{X}_k)$. Now using Boole's inequality (see Lemma 2.9) and

(5.16), we can write

$$\begin{aligned}\mathbb{P}\{\hat{\rho}_k < \eta_1\} &\leq \mathbb{P}\{|1-\hat{\rho}_k| \geq 1-\eta_1\} \\ &\leq \mathbb{P}\left\{\left|\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) - M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \geq (1-\eta_1)\,\kappa_{md}\Delta_k^2\right\} \\ &\leq \mathbb{P}\left\{Err_1 \geq \eta'\Delta_k^2\right\} + \mathbb{P}\left\{Err_2 \geq \eta'\Delta_k^2\right\} + \mathbb{P}\left\{Err_3 \geq \eta'\Delta_k^2\right\}, \end{aligned} \tag{5.17}$$

where $Err_1 := \left|\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) - f\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|$, $Err_2 := \left|f\left(\tilde{\boldsymbol{X}}_{k+1}\right) - m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|$, $Err_3 := \left|m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right) - M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right|$, and $\eta' = 3^{-1}(1-\eta_1)\,\kappa_{md}$. In what follows, we establish $\mathbb{P}\{\hat{\rho}_k < \eta_1 \text{ i.o.}\} = 0$ by demonstrating that each of the errors Err_1, Err_2 and Err_3 exceeding $\eta'\Delta_k^2$ infinitely often has probability zero.

We first analyze the stochastic sampling error probability $\mathbb{P}\left\{Err_1 \geq \eta'\Delta_k^2\right\}$ appearing on the right-hand side of (5.17). Using arguments identical to those leading to (5.6) in the proof of Lemma 5.1, it is seen that

$$\mathbb{P}\left\{\left|\bar{F}\left(\tilde{\boldsymbol{X}}_{k+1}, \widetilde{N}_{k+1}\right) - f\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \geq \eta'\Delta_k^2 \text{ i.o.}\right\} = 0. \tag{5.18}$$

Next we analyze the deterministic model error probability $\mathbb{P}\left\{Err_2 \geq \eta'\Delta_k^2\right\}$ appearing on the right-hand side of (5.17). Since we know from the postulates of the theorem that $m_k(\boldsymbol{z})$ is a $(\kappa_{ef}, \kappa_{eg})$-fully-linear model of f on $\mathcal{B}(\boldsymbol{X}_k; \Delta_k)$, implying that if η_1 is chosen so that $\eta' = \frac{1}{3}(1-\eta_1)\kappa_{md} > \kappa_{ef}$, we have

$$\mathbb{P}\left\{\left|f\left(\tilde{\boldsymbol{X}}_{k+1}\right) - m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \geq \eta'\Delta_k^2 \text{ i.o.}\right\} = 0. \tag{5.19}$$

Finally, we analyze the stochastic interpolation error probability $\mathbb{P}\left\{Err_3 \geq \eta'\Delta_k^2\right\}$ appearing on the right-hand side of (5.17). Using part (i) of Lemma 2.8 and relabeling $\boldsymbol{X}_k$ to $\boldsymbol{Y}_1$ for readability, we write

$$\begin{aligned}\mathbb{P}\left\{Err_3 > \eta'\Delta_k^2\right\} &\leq \mathbb{P}\left\{\max_{\substack{\boldsymbol{Y}_i \in \mathcal{Y}_k, \\ i=1,2,\ldots,p}} \left|\bar{F}\left(\boldsymbol{Y}_i, N\left(\boldsymbol{Y}_i\right)\right) - f\left(\boldsymbol{Y}_i\right)\right| > \frac{\eta'\Delta_k^2}{p\Lambda}\right\} \\ &\leq \sum_{i=1}^{p} \mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{Y}_i, N\left(\boldsymbol{Y}_i\right)\right) - f\left(\boldsymbol{Y}_i\right)\right| > \frac{\eta'\Delta_k^2}{p^2\Lambda}\right\} \\ &= \sum_{i=1}^{p} \mathbb{E}\left[\mathbb{P}\left\{\left|\bar{F}\left(\boldsymbol{Y}_i, N\left(\boldsymbol{Y}_i\right)\right) - f\left(\boldsymbol{Y}_i\right)\right| > \frac{\eta'\Delta_k^2}{p^2\Lambda} \,|\, \mathcal{F}_k\right\}\right]. \end{aligned} \tag{5.20}$$

Now using (5.20) and arguments identical to those leading to (5.6) in the proof of Lemma 5.1 (and the sample size expression (4.2) in Step 5 of Algorithm 2), we can then say for large enough k and some $\delta > 0$ that

$$\begin{aligned}\mathbb{P}\left\{Err_3 > \eta'\Delta_k^2\right\} &\leq \frac{p^5\Lambda}{\lambda_k(\eta'\Delta_k^2)^2}(1+\delta)\kappa_{ias}^2\Delta_k^4 \\ &\leq \frac{p^5\Lambda}{\lambda_k\eta'^2}(1+\delta)\kappa_{ias}^2. \end{aligned} \tag{5.21}$$

Since λ_k is chosen so that $k^{1+\epsilon} = \mathcal{O}(\lambda_k)$ for some $\epsilon > 0$, we see that (5.21) implies that $\mathbb{P}\left\{\left(\bar{F}\left(\boldsymbol{Y}_i, N\left(\boldsymbol{Y}_i\right)\right) - f\left(\boldsymbol{Y}_i\right)\right) > \eta'\Delta_k^2 \text{ i.o.}\right\} = 0$ by Borel-Cantelli. This in turn implies from (5.20) that

$$\mathbb{P}\left\{\left|M_k\left(\tilde{\boldsymbol{X}}_{k+1}\right) - m_k\left(\tilde{\boldsymbol{X}}_{k+1}\right)\right| \geq \eta'\Delta_k^2 \text{ i.o.}\right\} = 0. \tag{5.22}$$

Conclude from (5.18), (5.19), and (5.22) that each of the errors Err_1, Err_2, and Err_3 exceeding $\eta'\Delta_k^2$ infinitely often has probability zero and the assertion of Theorem 5.4 holds. □

LEMMA 5.5. *For any sample path $\omega \in \Omega$ if there exists a constant $\kappa_{lbg}(\omega) > 0$, such that $\|\nabla M_k(\boldsymbol{X}_k(\omega))\| \geq \kappa_{lbg}(\omega)$ for large enough k, then there exists a constant $\kappa_{lbd}(\omega) > 0$ such that $\Delta_k(\omega) \geq \kappa_{lbd}(\omega)$ for large enough k.*

Proof. Let $K_g(\omega) > 0$ be such that $\|\nabla M_k(\boldsymbol{X}_k(\omega))\| \geq \kappa_{lbg}$ if $k > K_g(\omega)$. From Theorem 5.4, we let $K_s(\omega) > 0$ be such that $K_s(\omega) - 1$ is the last unsuccessful iteration, that is, k is a successful iteration if $k \geq K_s(\omega)$. Then $\tilde{\Delta}_k(\omega) > \Delta_{k-1}(\omega)$ for all $k \geq K_s(\omega)$. For $k \geq \max\{K_g(\omega), K_s(\omega)\} + 1$, consider the two cases below when Algorithm 2 starts.

Case 1: $(\tilde{\Delta}_k(\omega) \geq \mu\|\nabla M_k(\boldsymbol{X}_k(\omega))\|)$: Since $\tilde{\Delta}_k(\omega) \geq \mu\|\nabla M_k(\boldsymbol{X}_k(\omega))\|$, the inner loop of Algorithm 2 is executed, implying that

$$\Delta_k(\omega) \geq \beta\|\nabla M_k(\boldsymbol{X}_k(\omega))\| \geq \beta\kappa_{lbg}(\omega).$$

Case 2: $(\tilde{\Delta}_k(\omega) < \mu\|\nabla M_k(\boldsymbol{X}_k(\omega))\|)$: In this scenario, the inner loop of Algorithm 2 is not executed, implying that $\Delta_k(\omega) = \tilde{\Delta}_k(\omega) = \gamma_1\Delta_{k-1}(\omega)$ meaning that the trust-region radius expands from the previous iteration.

Case 1 and Case 2 iterations are mutually exclusive and collectively exhaustive. Case 1 iterations imply, under the assumed postulates, that $\Delta_k(\omega) \geq \beta\kappa_{lbg}(\omega)$; Case 2 iterations result in an expanded trust-region radius. Conclude from these assertions that $\Delta_k(\omega) \geq \min\left\{\beta\kappa_{lbg}(\omega), \Delta_{\max\{K_g(\omega),K_s(\omega)\}}\right\}$. □

An important observation from the Algorithms 1 and 2 is that the difference between the function estimates of two consecutive iterates can be increasing; or in other words $\bar{F}(\boldsymbol{X}_k)$ is not necessarily monotone decreasing. When iteration k is unsuccessful, that is, $\boldsymbol{X}_k = \boldsymbol{X}_{k+1}$, it is possible that $\bar{F}(\boldsymbol{X}_k, N_k) < \bar{F}(\boldsymbol{X}_{k+1}, N_{k+1})$. When iteration k is successful, it must be true that $\bar{F}\left(\boldsymbol{X}_k, \widetilde{N}_k\right) > \bar{F}(\boldsymbol{X}_{k+1}, N_{k+1})$ but it is still possible that $\bar{F}(\boldsymbol{X}_k, N_k) < \bar{F}(\boldsymbol{X}_{k+1}, N_{k+1})$ since $\bar{F}(\boldsymbol{X}_k, N_k) \neq \bar{F}\left(\boldsymbol{X}_k, \widetilde{N}_k\right)$. We are now fully setup to demonstrate that ASTRO-DF's iterates converge to a first-order critical point with probability one.

THEOREM 5.6. *Let Assumptions 1 – 6 hold. Then $\|\nabla f(\boldsymbol{X}_k)\| \xrightarrow{wp1} 0$ as $k \to \infty$.*

Proof. We start by making the following two observations.

(a) Lemma 5.5 and Theorem 5.2 together imply that $\liminf_{k\to\infty}\|\nabla M_k(\boldsymbol{X}_k)\| = 0$. This, along with Lemma 5.3, imply $\liminf_{k\to\infty}\|\nabla f(\boldsymbol{X}_k)\| = 0$ almost surely.

(b) Consider the sequence of function estimates $\{\bar{F}(\boldsymbol{X}_k, N_k)\}$. Since we know from Theorem 5.4 that iterations are ultimately successful with probability one, we see that $\{\bar{F}(\boldsymbol{X}_k, N_k)\}$ is non-increasing for large enough k with probability one. This and the fact that the sequence $\{\bar{F}(\boldsymbol{X}_k, N_k)\}$ is bounded (by Lemma 5.1) with probability one implies that the difference $\bar{F}(\boldsymbol{X}_{k+1}, N_{k+1}) - \bar{F}(\boldsymbol{X}_k, N_k) \xrightarrow{wp1} 0$.

Suppose the assertion of Theorem 5.6 is not true. The there exists a set $\mathcal{D}$ of positive measure such that for any sample-path $\omega \in \mathcal{D}$, there exists a subsequence of iterations $\{t_i\}$ satisfiying $\|\nabla f(\boldsymbol{X}_{t_i})\| > 3\epsilon$ for some $\epsilon > 0$. (In the previous statement and in what follows, we have suppressed ω from the notation for convenience.) Due to our observation in (a), corresponding to each element t_i, there exists $\ell_i = \ell(t_i)$, the first iteration after t_i, such that $\|\nabla f(\boldsymbol{X}_{\ell_i})\| < 2\epsilon$. Therefore if $\mathcal{K}_i = \{k : t_i \leq k \leq \ell_i\}$, then $\|\nabla f(\boldsymbol{X}_k)\| \geq 2\epsilon$ for all $k \in \mathcal{K}_i$.

Now choose i large enough so that by Theorem 5.4, Lemma 5.3, Theorem 5.2, and Lemma 5.1 for all $k \in \mathcal{K}_i$, (i) $\hat{\rho}_k \geq \eta_1$ (only successful iterations), (ii) $\|\nabla M_k(\boldsymbol{X}_k)\| \geq \epsilon$ (model gradient close to the function gradient), (iii) $\Delta_k \leq \kappa_{bhm}^{-1}\epsilon$ (trust-region radius small), and (iv) $\left|\bar{F}(\boldsymbol{X}_k, N_k) - f(\boldsymbol{X}_k)\right| \leq 8^{-1}\eta_1\epsilon\kappa_{fcd}\Delta_k$ (simulation error small). As a result and by Cauchy reduction in Assumption 4 we get

$$\begin{aligned}\bar{F}(\boldsymbol{X}_{k+1}, N_{k+1}) - \bar{F}\left(\boldsymbol{X}_k, \widetilde{N}_k\right) &\leq -2^{-1}\eta_1\kappa_{fcd}\|\nabla M_k(\boldsymbol{X}_k)\| \min\left\{\kappa_{bhm}^{-1}\|\nabla M_k(\boldsymbol{X}_k)\|, \Delta_k\right\} \\ &\leq -2^{-1}\eta_1\kappa_{fcd}\epsilon \min\left\{\kappa_{bhm}^{-1}\epsilon, \Delta_k\right\} \\ &= -\eta_1\frac{\kappa_{fcd}}{2}\epsilon\Delta_k. \qquad (5.23)\end{aligned}$$

Therefore

$$\begin{aligned}\bar{F}(\boldsymbol{X}_{k+1}, N_{k+1}) - \bar{F}(\boldsymbol{X}_k, N_k) &\leq \bar{F}(\boldsymbol{X}_{k+1}, N_{k+1}) - \bar{F}\left(\boldsymbol{X}_k, \widetilde{N}_k\right) \\ &\quad + \left|\bar{F}\left(\boldsymbol{X}_k, \widetilde{N}_k\right) - f(\boldsymbol{X}_k)\right| + \left|f(\boldsymbol{X}_k) - \bar{F}(\boldsymbol{X}_k, N_k)\right| \\ &\leq -\eta_1\frac{\kappa_{fcd}}{2}\epsilon\Delta_k + \eta_1\frac{\kappa_{fcd}}{4}\epsilon\Delta_k = -\eta_1\frac{\kappa_{fcd}}{4}\epsilon\Delta_k, \qquad (5.24)\end{aligned}$$

and as a result $\Delta_k \leq -4(\eta_1\kappa_{fcd}\epsilon)^{-1}\left(\bar{F}(\boldsymbol{X}_{k+1}, N_{k+1}) - \bar{F}(\boldsymbol{X}_k, N_k)\right)$ for all $k \in \mathcal{K}_i$. It follows that

$$\begin{aligned}\|\boldsymbol{X}_{\ell_i} - \boldsymbol{X}_{t_i}\| &\leq \sum_{j\in\mathcal{K}_i}\|\boldsymbol{X}_{j+1} - \boldsymbol{X}_j\| \leq \sum_{j\in\mathcal{K}_i}\Delta_j \\ &\leq \frac{-4}{\eta_1\kappa_{fcd}\epsilon}\sum_{j\in\mathcal{K}_i}\bar{F}(\boldsymbol{X}_{j+1}, N_{j+1}) - \bar{F}(\boldsymbol{X}_j, N_j) \\ &\leq \frac{-4}{\eta_1\kappa_{fcd}\epsilon}\left(\bar{F}(\boldsymbol{X}_{\ell_i}, N_{\ell_i}) - \bar{F}(\boldsymbol{X}_{t_i}, N_{t_i})\right). \qquad (5.25)\end{aligned}$$

The inequality in (5.25) and our observation in (b) imply that

$$\|\boldsymbol{X}_{\ell_i} - \boldsymbol{X}_{t_i}\| \to 0 \text{ as } i \to \infty. \qquad (5.26)$$

Furthermore, since

$$\begin{aligned}|f(\boldsymbol{X}_{\ell_i}) - f(\boldsymbol{X}_{t_i})| &\leq \left|\bar{F}(\boldsymbol{X}_{\ell_i}, N_{\ell_i}) - \bar{F}(\boldsymbol{X}_{t_i}, N_{t_i})\right| \\ &\quad + \left|f(\boldsymbol{X}_{\ell_i}) - \bar{F}(\boldsymbol{X}_{\ell_i}, N_{\ell_i})\right| + \left|f(\boldsymbol{X}_{t_i}) - \bar{F}(\boldsymbol{X}_{t_i}, N_{t_i})\right|,\end{aligned}$$

we see that

$$|f(\boldsymbol{X}_{\ell_i}) - f(\boldsymbol{X}_{t_i})| \to 0 \text{ as } i \to \infty. \qquad (5.27)$$

Using (5.26) and (5.27), and since the gradient $\nabla f(\boldsymbol{x})$ is Lipschitz continuous, we conclude $\|\nabla f(\boldsymbol{X}_{\ell_i}) - \nabla f(\boldsymbol{X}_{t_i})\| \to 0$ as $i \to \infty$. This, however, gives us a contradiction since the definition of t_i and ℓ_i dictate that $\|\nabla f(\boldsymbol{X}_{t_i})\| > 3\epsilon$ and $\|\nabla f(\boldsymbol{X}_{\ell_i})\| < 2\epsilon$.□

6. FURTHER REMARKS AND DISCUSSION. Over the last decade or so, derivative-free trust-region algorithms have deservedly enjoyed great attention and success in the deterministic optimization context. Analogous algorithms for the now widely prevalent and important Monte Carlo stochastic optimization context,

where only stochastic function oracles are available, is poorly studied. This paper develops adaptive sampling trust-region optimization derivative-free algorithms (called ASTRO-DF) for solving low to moderate dimensional stochastic optimization problems. The key idea within ASTRO-DF is to endow a derivative-free trust-region algorithm with an adaptive sampling strategy for function estimation. The extent of such sampling at a visited point depends on the estimated proximity of the point to a solution, calculated by balancing the estimated standard error of the function estimate with a certain power of the incumbent trust-region radius. So, just as one might expect of efficient algorithms, Monte Carlo sampling in ASTRO-DF tends to be low during the early iterations compared to later iterations, when the visited points are more likely to be closer to a first-order critical point. More importantly, however, the schedule of sampling is not predetermined (as in most stochastic approximation and sample-average approximation algorithms) but instead adapts to the prevailing algorithm trajectory and the needed precision of the function estimates.

We show that ASTRO-DF's iterates exhibit global convergence to a first-order critical point with probability one. While the proofs are detailed, convergence follows from two key features of ASTRO-DF: (i) the stochastic interpolation models are constructed across iterates in such a way that the error in the stochastic interpolation model is guaranteed to remain in lock-step with (a certain power of) the trust-region radius; and (ii) the optimization within the trust-region step is performed in such a way as to guarantee Cauchy reduction and then the objective function is evaluated at the resulting candidate point. Remarkably, the features (i) and (ii) together ensure that the sequence of trust-region radii necessarily need to converge to zero with probability one, and that the model gradient, the true gradient and the trust-region radius all have to remain in lock-step, thus guaranteeing convergence to a first-order critical point with probability one. The key driver for efficiency is adaptive sampling, making all sample sizes within ASTRO-DF stopping times that are explicitly dependent on algorithm trajectory.

Four other points are worthy of mention.

(i) Our proofs demonstrate global convergence to first-order critical points. Corresponding proofs of convergence to a second-order critical point can be obtained in an identical fashion by driving some measure of second-order stationarity to zero instead of the model gradient.

(ii) The adaptive sampling ideas and the ensuing proofs we have presented in this paper are for the specific case of stochastic interpolation models. It seems to us, however, that the methods of proof presented in this paper can be co-opted (with care) into other potentially more powerful model construction ideas such as regression [14] and kriging [3, 56]. Which of such ideas result in the best derivative-free trust-region algorithms (as measured by practical performance and asymptotic efficiency) remains to be seen.

(iii) We have presented no proof that ASTRO-DF's iterates achieve the Monte Carlo canonical rate [4]. Demonstrating that ASTRO-DF's iterates achieve the canonical rate will rely on rate results for derivative-free trust-region algorithms in the deterministic context, some of which are only now appearing [26]. We speculate, however, that ASTRO-DF's iterates do enjoy the canonical rate, as our analogous work [29] in a different context has demonstrated.

(iv) The asymptotic sampling rate within ASTRO-DF is approximately $\mathcal{O}\left(\Delta_k^{-4}\right)$, where Δ_k is the incumbent trust-region radius. (See (5.4) in the proof of

Lemma 5.1.) This sampling stipulation is comparable to that prescribed in two other prominent recent studies [19, 36].The $\mathcal{O}(\Delta^{-4})$ sampling appears to be the minimum needed to guarantee convergence in derivative-free trust-region methods *without assumptions on the tail-behavior of the error driving simulation observations.* A question of interest is whether the $\mathcal{O}(\Delta^{-4})$ sampling stipulation can be relaxed by assuming that the simulation error is light-tailed, that is, the errors have a well-defined moment-generating function. Such assumption will likely allow modifying the results in [19] through the use of the Chernoff bound instead of the Chebyshev inequality, leading to a weakening of the sampling stipulation. The corresponding modification in ASTRO-DF will involve proving a variation of Theorem 3.7 which we believe will be a contribution in itself. Assuming that the simulation errors are light-tailed is reasonable. For instance, many distributions we see in practice, e.g, normal, gamma, beta, are light-tailed; any distribution with bounded support is light-tailed.

Appendix A. Proof of part (i) of Lemma 2.8. We know that for all $\boldsymbol{z} \in \mathcal{B}(\boldsymbol{Y}_1; \Delta)$,

$$m(\boldsymbol{z}) = \sum_{i=1}^{p} \ell_i(\boldsymbol{z}) f(\boldsymbol{Y}_i); \; M(\boldsymbol{z}) = \sum_{i=1}^{p} \ell_i(\boldsymbol{z}) \bar{F}(\boldsymbol{Y}_i, n(\boldsymbol{Y}_i)),$$

where $\ell_j(\boldsymbol{z})$ are the Lagrange polynomials associated with the set $\mathcal{Y}$. Since $\mathcal{Y}$ is Λ-poised in $\mathcal{B}(\boldsymbol{Y}_1; \Delta)$, we know (see Chapter 3 in [22]) that

$$\Lambda \geq \Lambda_\ell = \max_{i=1,2,\ldots,p} \max_{\boldsymbol{z} \in \mathcal{B}(\boldsymbol{Y}_1;\Delta)} |\ell_i(\boldsymbol{z})|. \tag{A.1}$$

Now write, for $\boldsymbol{z} \in \mathcal{B}(\boldsymbol{Y}_1; \Delta)$,

$$\begin{aligned} |M(\boldsymbol{z}) - m(\boldsymbol{z})| &= \left| \sum_{i=1}^{p} \ell_i(\boldsymbol{z}) \left(\bar{F}(\boldsymbol{Y}_i, n(\boldsymbol{Y}_i)) - f(\boldsymbol{Y}_i) \right) \right| \\ &\leq p\Lambda_\ell \max_{i \in \{1,2,\ldots,p\}} \left| \bar{F}(\boldsymbol{Y}_i, n(\boldsymbol{Y}_i)) - f(\boldsymbol{Y}_i) \right| \\ &\leq p\Lambda \max_{i \in \{1,2,\ldots,p\}} \left| \bar{F}(\boldsymbol{Y}_i, n(\boldsymbol{Y}_i)) - f(\boldsymbol{Y}_i) \right|, \end{aligned}$$

where the last inequality follows from (A.1).

Appendix B. Proof of part (ii) of Lemma 2.8.

If the model $M(\cdot)$ is a stochastic linear interpolation model, we see that for $i = 1, 2, 3, \ldots, p$

$$(\boldsymbol{Y}_i - \boldsymbol{Y}_1)^T \nabla M(\boldsymbol{Y}_1) = M(\boldsymbol{Y}_i) - M(\boldsymbol{Y}_1) = f(\boldsymbol{Y}_i) - f(\boldsymbol{Y}_1) + E_i - E_1. \tag{B.1}$$

Now re-trace the steps of the proof on pages 26 and 27 of [22], while carrying the additional term $E_i - E_1$ appearing on the right-hand side of (B.1).

If the model $M(\cdot)$ is a stochastic quadratic interpolation model, we write $\boldsymbol{z} \in \mathcal{B}(\boldsymbol{Y}_1, \Delta)$, $M(\boldsymbol{z}) = c + \boldsymbol{z}^T g + \frac{1}{2}\boldsymbol{z}^T H \boldsymbol{z} = f(\boldsymbol{z}) + e^f(\boldsymbol{z}); \nabla M(\boldsymbol{z}) = H\boldsymbol{z} + g = \nabla f(\boldsymbol{z}) + e^g(\boldsymbol{z}); \nabla^2 M(\boldsymbol{z}) = H = \nabla^2 f(\boldsymbol{z}) + e^H(\boldsymbol{z})$. Now write, after subtracting the expression for $M(\boldsymbol{z})$ from that for $M(\boldsymbol{Y}_i), i = 1, \ldots, p+1$ to get

$$(\boldsymbol{Y}_i - \boldsymbol{z})^T g + \frac{1}{2}(\boldsymbol{Y}_i - \boldsymbol{z})^T H (\boldsymbol{Y}_i - \boldsymbol{z}) + (\boldsymbol{Y}_i - \boldsymbol{z})^T H \boldsymbol{z} = f(\boldsymbol{Y}_i) - f(\boldsymbol{z}) + E^i - e^f(\boldsymbol{z}). \tag{B.2}$$

Notice that the equation in (B.2) is identical to the corresponding equation on page 53 of [22] except for the extra term E^i appearing on the right-hand side of (B.2). Re-trace the steps on pages 53, 54, and 55 of [22].

Appendix C. Model Construction Algorithm Termination. In what follows, we demonstrate through the following result that the model construction algorithm (Algorithm 2) terminates with probability one, whenever the incumbent solution $\boldsymbol{X}_k$ is not a first-order critical point.

LEMMA C.1. *Suppose the incumbent solution $\boldsymbol{X}_k \in \mathbb{R}^d$ during the kth iteration is not first-order critical, that is, $\nabla f(\boldsymbol{X}_k) \neq 0$. Then Algorithm 2 terminates in a finite number of steps with probability one.*

Proof. Set $\|\nabla f(\boldsymbol{X}_k)\| = c' \neq 0$. We will prove the assertion through a contradiction argument.

First, we notice that the contraction loop (Steps 3–9) in Algorithm 2 is not entered if $\mu \|\nabla M_k(\boldsymbol{X}_k)\| \geq \tilde{\Delta}_k$, in which case Algorithm 2 terminates trivially.

Next, suppose $\mu \|\nabla M_k(\boldsymbol{X}_k)\| < \tilde{\Delta}_k$ and that the contraction loop in Steps 3–9 of Algorithm 2 is infinite. Let $\nabla M_k^{(j_k)}(\boldsymbol{X}_k)$ denote the model gradient during the k_jth iteration of the contraction loop. Then $\mu \left\|\nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\| < \tilde{\Delta}_k w^{j_k-1}, \ \forall j_k \geq 1$. This means, since $w < 1$, that $\tilde{\Delta}_k w^{j_k-1} \to 0$ and therefore $\left\|\nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\| \xrightarrow{wp1} 0$ as $j_k \to \infty$. Furthermore, due to the sampling rule in (4.2) and by Theorem 2.6, we have that $N\left(\boldsymbol{Y}_i^{(j_k)}\right) \to \infty$ as $j_k \to \infty$. Now, if $E_{k,i}^{(j_k)} = \bar{F}\left(\boldsymbol{Y}_i^{(j_k)}, N\left(\boldsymbol{Y}_i^{(j_k)}\right)\right) - f\left(\boldsymbol{Y}_i^{(j_k)}\right)$, then we can write for large enough k and some $\delta > 0$,

$$\mathbb{P}\left\{\frac{\sum_{i=2}^{p}\left|E_{k,i}^{(j_k)} - E_{k,1}^{(j_k)}\right|}{\tilde{\Delta}_k w^{j_k-1}} \geq c\right\} \leq 8(p-1)^3 c^{-2}(1+\delta)\kappa_{ias}^2 \Delta_k^2 \lambda_k^{-1}, \tag{C.1}$$

which follows from arguments identical to (5.15) in the proof of Lemma 5.3. Since the right-hand side of (C.1) is summable, we conclude by Borel-Cantelli's first lemma (Lemma 2.10) that $\left(\tilde{\Delta}_k w^{j_k-1}\right)^{-1} \sum_{i=2}^{p}\left|E_{k,i}^{(j_k)} - E_{k,1}^{(j_k)}\right| \xrightarrow{wp1} 0$. This implies, from Lemma 2.8 and since Algorithm 2 maintains full-linearity, that as $j_k \to \infty$,

$$\left\|\nabla f(\boldsymbol{X}_k) - \nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\| \leq \kappa_1\left(\tilde{\Delta}_k w^{j_k-1}\right)^{\theta} + \kappa_2 \frac{\sum_{i=1}^{p}\left|E_{k,i}^{(j_k)} - E_{k,1}^{(j_k)}\right|}{\left(\tilde{\Delta}_k w^{j_k-1}\right)} \xrightarrow{wp1} 0.$$

Hence we have arrived at a contradiction since we argued that $\left\|\nabla M_k^{(j_k)}(\boldsymbol{X}_k)\right\| \xrightarrow{wp1} 0$ but then $\|\nabla f(\boldsymbol{X}_k)\| = c' \neq 0$ by the contrapositive assumption. □

REFERENCES

[1] O. Alagoz, A. J. Schaefer, and M. S. Roberts. Optimization in Organ Allocation. In P. Pardalos and E. Romeijn, editors, *Handbook of Optimization in Medicine*. Kluwer Academic Publishers, 2009.

[2] B. D. Amos, D. R. Easterling, L. T. Watson, W. I. Thacker, B. S. Castle, and M. W. Trosset. Algorithm XXX: QNSTOP Quasi Newton Algorithm for Stochastic Optimization. 2014.

[3] B. Ankenman, B. L. Nelson, and J. Staum. Stochastic Kriging for Simulation Metamodeling. *Operations research*, 58(2):371–382, 2010.

[4] S. Asmussen and P. W. Glynn. *Stochastic Simulation: Algorithms and Analysis*, volume 57. Springer Science & Business Media, 2007.

[5] A. S. Bandeira, K. Scheinberg, and L. N. Vicente. Convergence of Trust-Region Methods Based on Probabilistic Models. *SIAM Journal on Optimization*, 24(3):1238–1264, 2014.

[6] F. Bastin, C. Cirillo, and P. L. Toint. Convergence Theory for Nonconvex Stochastic Programming with an Application to Mixed Logit. *Mathematical Programming*, 108:207–234, 2006.

[7] F. Bastin, C. Cirillo, and Ph. L. Toint. An Adaptive Monte Carlo Algorithm for Computing Mixed Logit Estimators. *Computational Management Science*, 3(1):55–79, 2006.

[8] G. Bayraksan and D. P. Morton. Assessing Solution Quality in Stochastic Programs. *Mathematical Programming Series B*, 108:495–514, 2007.

[9] G. Bayraksan and D. P. Morton. A Sequential Sampling Procedure for Stochastic Programming. *Operations Research*, 59(4):898–913, 2011.

[10] G. Bayraksan and P. Pierre-Louis. Fixed-Width Sequential Stopping Rules for a Class of Stochastic Programs. *SIAM Journal on Optimization*, 22(4):1518–1548, 2012.

[11] M. S. Bazaara, H. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. John Wiley & Sons, New York, NY., 2006.

[12] P. Billingsley. *Probability and Measure*. Wiley, New York, NY., 1995.

[13] S. C. Billups, J. Larson, and P. Graf. Derivative-Free Optimization of Expensive Functions with Computational Error Using Weighted Regression. *SIAM Journal on Optimization*, 23(1):27–53, 2013.

[14] S. C. Billups, J. Larson, and P. Graf. Derivative-Free Optimization of Expensive Functions with Computational Error Using Weighted Regression. *SIAM Journal on Optimization*, 23(1):27–53, 2013.

[15] J. Blum. Approximation methods which converge with probability one. *Annals of Mathematical Statistics*, 25(2):382–386, 1954.

[16] M. Broadie, D. M. Cicek, and A. Zeevi. An Adaptive Multidimensional Version of the Kiefer-Wolfowitz Stochastic Approximation Algorithm. In M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, editors, *Proceedings of the 2009 Winter Simulation Conference*, pages 601–612. Institute of Electrical and Electronics Engineers: Piscataway, New Jersey, 2009.

[17] M. Broadie, D. M. Cicek, and A. Zeevi. General Bounds and Finite-Time Improvement for the Kiefer-Wolfowitz Stochastic Approximation Algorithm. *Operations Research*, 59(5):1211–1224, 2011.

[18] K. Chang, L. J. Hong, and H. Wan. Stochastic Trust-Region Response-Surface Method (STRONG) - a New Response-Surface Framework for Simulation Optimization. *INFORMS Journal on Computing*, 25(2):230–243, 2013.

[19] R. Chen, M. Menickelly, and K. Scheinberg. Stochastic Optimization Using a Trust-Region Method and Random Models, 2015. submitted.

[20] T. S. Chow and H. Robbins. On the Asymptotic Theory of Fixed-Width Sequential Confidence Intervals for the Mean. *The Annals of Mathematical Statistics*, pages 457–462, 1965.

[21] A. R. Conn, K. Scheinberg, and L. N. Vicente. Global Convergence of General Derivative-Free trust-Region Algorithms to First-and Second-Order Critical Points. *SIAM Journal on Optimization*, 20(1):387–415, 2009.

[22] A. R. Conn, K. Scheinberg, and L. N. Vicente. *Introduction to Derivative-Free Optimization*, volume 8. Siam, 2009.

[23] G. Deng and M. C. Ferris. Adaptation of the UOBYQA Algorithm for Noisy Functions. In *Proceedings of the 38th conference on Winter simulation*, pages 312–319. Winter Simulation Conference, 2006.

[24] G. Deng and M. C. Ferris. Variable-Number Sample-Path Optimization. *Mathematical Programming*, 117(1-2):81–109, 2009.

[25] R. Durrett. *Probability: Theory and Examples*. Cambridge University Press, New York, NY, 2010.

[26] R. Garmanjani, D. Júdice, and L. N. Vicente. Trust-Region Methods Without Using Derivatives: Worst Case Complexity and the Non-Smooth Case. 2015.

[27] M. Ghosh and N. Mukhopadhyay. Sequential Point Estimation of the Mean When the Distribution is Unspecified. *Communications in Statistics - Theory and Methods*, pages 637–652, 1979.

[28] M. Ghosh, N. Mukhopadhyay, and P. K. Sen. *Sequential Estimation*. Wiley Series in Probability and Statistics, 1997.

[29] F. S. Hashemi, S. Ghosh, and R. Pasupathy. On Adaptive Sampling Rules for Stochastic Recursions. In *Proceedings of the 2014 Winter Simulation Conference*, pages 3959–3970.

IEEE Press, 2014.
[30] S. G. Henderson and B. L. Nelson, editors. volume 13 of *Handbooks in Operations Research and Management Science: Simulation*. Elsevier, 2006.
[31] Y. T. Hou, Y. Shi, and H. D. Sherali. *Applied Optimization Methods for Wireless Networks*. Cambridge University Press, 2014.
[32] J. Kiefer and J. Wolfowitz. Stochastic Estimation of the Maximum of a Regression Function. *Annals of Mathematical Statistics*, 23:462–466, 1952.
[33] S. Kim, R. Pasupathy, and S. G. Henderson. A Guide to SAA. Frederick Hilliers OR Series. Elsevier, 2014.
[34] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, New York, NY., 2003.
[35] J. Larson and S. C. Billups. Stochastic Derivative-Free Optimization Using a Trust-Region Framework. *Under review at Computational Optimization and Applications*, 2014.
[36] Jeffrey M. Larson. *Derivative-Free Optimization of Noisy Functions*. PhD thesis, Department of Applied Mathematics, University of Colorado, Denver, CO, 2012.
[37] W. K. Mak, D. P. Morton, and R. K. Wood. Monte Carlo Bounding Techniques for Determining Solution Quality in Stochastic Programs. *Operations Research Letters*, 24:47–56, 1999.
[38] A. Mokkadem and M. Pelletier. A Generalization of the Averaging Procedure: The Use of Two-Time-Scale Algorithms. *SIAM Journal on Control and Optimization*, 49:1523, 2011.
[39] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer-Verlag, Berlin, 2006.
[40] E. O. Nsoesie, R. J. Beckman, S. Shashaani, K. S. Nagaraj, and M. V. Marathe. A Simulation Optimization Approach to Epidemic Forecasting. *PloS one*, 8(6):e67164, 2013.
[41] J. M. Ortega and W. C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. Academic Press, New York, NY., 1970.
[42] C. Osorio and M. Bierlaire. A Surrogate Model for Traffic Optimization of Congested Networks: an Analytic Queueing Network Approach. *Report TRANSP-OR*, 90825:1–23, 2009.
[43] R. Pasupathy. On Choosing Parameters in Retrospective-Approximation Algorithms for Stochastic Root Finding and Simulation Optimization. *Operations Research*, 58:889–901, 2010.
[44] R. Pasupathy and S. Ghosh. Simulation Optimization: A Concise Overview and Implementation Guide. INFORMS TutORials. INFORMS, 2013.
[45] R. Pasupathy, P. W. Glynn, S. G. Ghosh, and F. S. Hashemi. How Much to Sample in Simulation-Based Stochastic Recursions? 2014. Under Review.
[46] R. Pasupathy and S. G. Henderson. A Testbed of Simulation-Optimization Problems. In *Proceedings of the 38th conference on Winter simulation*, pages 255–263. Winter Simulation Conference, 2006.
[47] R. Pasupathy and S. G. Henderson. SimOpt: A Library of Simulation Optimization Problems. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, editors, *Proceedings of the 2011 Winter Simulation Conference*. Institute of Electrical and Electronics Engineers: Piscataway, New Jersey, 2011.
[48] B. T. Polyak and A. B. Juditsky. Acceleration of Stochastic Approximation by Averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
[49] M. J. D. Powell. UOBYQA: Unconstrained Optimization by Quadratic Approximation. *Mathematical Programming*, 92(3):555–582, 2002.
[50] H. Robbins and S. Monro. A Stochastic Approximation Method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
[51] J. Royset and R. Szechtman. Optimal Budget Allocation for Sample Average Approximation. *Operations Research*, 2011. Under Review.
[52] A. Ruszczynski and A. Shapiro, editors. *Stochastic Programming. Handbook in Operations Research and Management Science*. Elsevier, New York, NY., 2003.
[53] A. Shapiro, D. Dentcheva, and A. Ruszczynski. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia, PA, 2009.
[54] J. C. Spall. Adaptive Stochastic Approximation by the Simultaneous Perturbation Method. *IEEE Transactions on Automatic Control*, 45:1839–1853, 2000.
[55] J. C. Spall. *Introduction to Stochastic Search and Optimization*. John Wiley & Sons, Inc., Hoboken, NJ., 2003.
[56] M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, New York, NY, 1999.
[57] F. Yousefian, A Nedić, U. V., and Shanbhag. On Stochastic Gradient and Subgradient Methods with Adaptive Step Length Sequences. *Automatica*, 48(1):56–67, 2012.